

Institute for Policy Integrity

New York University School of Law

January 13, 2010

VIA ELECTRONIC SUBMISSION

Federal Communications Commission
Attention: GN Docket No. 09–191, WC Docket No. 07–52, FCC 09–93

Subject: Comments on Proposed Rulemaking, “Preserving the Open Internet, Broadband Industry Practices,” 74 Fed. Reg. 62,638 (Nov. 30, 2009), GN Docket No. 09–191, WC Docket No. 07–52, FCC 09–93

The Institute for Policy Integrity submits the following comments in support of the efforts of the Federal Communications Commission (FCC) to address network neutrality on the Internet. The Institute for Policy Integrity (IPI) at New York University School of Law is a non-partisan advocacy organization and think-tank dedicated to improving the quality of government decisionmaking in the areas of environmental, public health, and consumer regulation. IPI advocates using rational economic analysis as a tool to advance socially-beneficial regulation.

Over the course of 2009, IPI conducted an economic analysis demonstrating that the FCC’s proposed rules concerning net neutrality would likely be economically beneficial to the U.S. Internet market. This Report, entitled *Free to Invest: The Economic Benefits of Preserving Net Neutrality*, is attached to these comments. It analyzes federal net neutrality policy from an economic perspective in order to understand the fundamental tradeoffs involved in Internet policy, and determines how to maximize the value of the Internet. While there may be other considerations for policymakers, economic criteria can indicate which policy maximizes net benefits for society and, thus, are clearly important.

The Report makes five core findings integral to any net neutrality policy:

Internet Market Failure

The Internet—understood both as the physical infrastructure as well as the content and information moving along that infrastructure—produces billions of dollars of free value for the American public: Information is shared, reused, and reconfigured without fees or penalties. Websites are not compensated when their content is repurposed or passed on—that means fewer subscriptions to paid services, fewer direct page views, and a loss of advertising dollars. This economic dynamic has been taken for granted as the Internet grew around the idea that information resides in the public domain—free to be emailed, Tweeted, blogged, and discussed.

Smart Policy Can Help

As a result of this dynamic, the Internet is more useful to everyone on it, but Internet Service Providers (ISPs) and content providers are at a disadvantage since they are not compensated for all the information they disseminate. This leads to systematic underinvestment in the Internet: if that income could be accessed, it would encourage investment in infrastructure and content. Government policy can increase investment in both Internet content and infrastructure to overcome this market failure.

Transferring Wealth Through Price Discrimination

Without net neutrality rules, new technologies could lead to pricing practices that transfer wealth from content providers to ISPs, a form of price discrimination that would reduce the return on investment for Internet content—meaning website owners, bloggers, newspapers, and businesses would have less incentive to expand their sites and applications.

Efficiently Supporting Infrastructure

Additional investment in broadband infrastructure would also increase the value of the Internet—making it faster and accessible in more places. But charging content providers for access to ISP customers is an extremely inefficient economic tool to generate such investment, primarily because most additional revenue generated for ISPs is likely to be transferred to their shareholders rather than invested in expanding broadband lines.

It is relatively easy to directly support infrastructure development, but hard to provide direct support for content. Targeted government support for ISPs to expand access where needed, along with net neutrality rules to protect content providers, are the best combination of policies for overcoming the market failure of underinvestment in the Internet.

Problems With Prioritization

Without net neutrality rules preventing priority pricing techniques, there could be changes in the way content appears online. If ISPs create “priority” or “fast lane” access to content providers at a fee, users could experience uneven access to websites and applications. While some content providers may benefit from this architecture, many type of websites will be especially harmed. Ultimately, prioritization could reduce incentives for content creators, potentially lowering the overall value of the Internet for all users.

By giving players the best incentives for optimal investment, net neutrality encourages a cycle that breeds more content, which in turn breeds more users. A combination of policies that protect content providers and judiciously deploy government resources to augment private investment in physical infrastructure is the right mix to ensure that the Internet continues to grow and flourish, generating massive benefits for the American public.

Conclusion

We applaud the agency’s efforts to craft rules that maximize the value of the Internet for American society. The recommendations contained in these comments will allow the FCC to more efficiently and effectively meet these goals.

Respectfully submitted,

Michael A. Livermore (mlivermore@nyu.edu)
Inimai M. Chettiar (inimai.chettiar@nyu.edu)
INSTITUTE FOR POLICY INTEGRITY AT
NEW YORK UNIVERSITY SCHOOL OF LAW

Institute for Policy Integrity
New York University School of Law



Free to Invest

The Economic Benefits of
Preserving Net Neutrality

Inimai M. Chettiar
J. Scott Holladay
Report No. 4
January 2010

Institute *for* Policy Integrity
New York University School of Law

Free to Invest

The Economic Benefits of Preserving Net Neutrality

Inimai M. Chettiar
J. Scott Holladay

Report No. 4
January 2010

Copyright © 2010 by the Institute for Policy Integrity.
All rights reserved.

Printed in the United States of America.

Institute for Policy Integrity
New York University School of Law
245 Sullivan Street, Suite 472
New York, New York 10012

Contents

Acknowledgements.....	iv
Foreword.....	v
Executive Summary.....	vii

Introduction: The Internet Policy Landscape	1
I. Positive Externalities: The Hidden Value of the Internet	7
II. The Net Neutrality Wealth Tradeoff: Content Providers and Internet Service Providers	17
III. The Problem with Price Discrimination.....	24
IV. Improving Infrastructure Through Direct Government Support	33
V. Fast Lanes to Nowhere: The Risks of Prioritization	39
Conclusion: Toward a Smart Internet Policy.....	48

Acknowledgments

This Report benefited from the invaluable guidance of Richard L. Revesz and Michael A. Livermore, and it expands upon the ideas expressed in their book, *Retaking Rationality: How Cost-Benefit Analysis Can Better Protect the Environment and Our Health* (Oxford University Press, 2008).

Thanks are also due to Edna Ishayik and Josip Markus, as well as to our Research Assistants Stephen Mar, Vivek Chandrasekhar, and James Thompson.

Foreword

The primary goals of the FCC's Internet policy should to ensure that the growth of the Internet continues, and to protect the innovation, creativity, and sharing that the Internet empowers.

The Internet has dominated the past decade. It has drastically altered the daily lives of millions of people: families separated by continents communicate on their computer screens as if they are in the same room; shoppers buy clothes, music, cars, and homes online; and experts share information about everything from cupcakes to particle physics.

Across the globe, so many of the defining political and cultural moments of the last ten years centered around the electronic networks that connect us: political protests organized from smart phones; news of natural disasters breaking online and spreading virally; and a presidential debate becoming a nationwide town hall as voters sent in questions from their computers.

The way we have come to use the Internet is a function of its openness—the cost of starting a website and sharing content with the world is low. Anyone with a few hundred dollars can buy a domain name, rent space on a server, and begin publishing content that anyone with an Internet connection can access. People with new ideas are encouraged to test them out as the number of users online can make the pay-off well worth the investment.

For every YouTube, Wikipedia, or Google there are thousands of websites and applications created and tested—some are game-changers, and some are not. But the depth and breadth of content is what drives the Internet to become wider, smarter, and more useful as each day passes.

Behind the scenes, this dynamic, referred to as “network neutrality” (or “net neutrality”) works like this: end-users pay Internet Service Providers (ISPs), like Verizon or Comcast, for access to the Internet; and content providers, like newspapers, blogs, and businesses, only pay ISPs fees to upload their information onto the Internet. Without net neutrality, ISPs could charge content providers again when users *download* that content. Adding these fees would increase the costs of creating websites and applications.

In October 2009, the Federal Communications Commission (FCC) proposed a set of net neutrality rules that would require ISPs to give equal access to all providers of Internet content. The proposed rules respond to new techniques that would enable ISPs to discriminate against certain types of

content in the provision of network services—for example, by giving certain applications “priority” when lines are congested.

Opponents of the proposed rules argue that they will stifle innovation and unfairly shackle ISPs from enjoying fair returns on their investment in Internet infrastructure. As the Internet has grown in recent decades, it has continually evolved in new directions that were first unimaginable, then novel, and then quickly became commonplace. Some believe that the Internet is best left alone, and fear that net neutrality rules will dampen the creative spirit at the heart of the successful growth of the Internet.

Proponents of net neutrality argue that the rules are needed to preserve innovation and protect content providers from being taken advantage of by ISPs, who are the gatekeepers for Internet use across the country. The next new technologies or Internet applications often do not come from established players; it is the innovators and start-up companies that provide the creative firepower that drives the evolution of the Internet forward. If ISP pricing crowds out this innovation, or significantly diminishes the returns for those who invest in creating new and novel forms of content, then the Internet will stagnate.

In this report, Inimai M. Chettiar and J. Scott Holladay discuss the implications of the proposed new net neutrality rules for the Internet, examining arguments on both side of the debate. Using an economic framework, Chettiar and Holladay analyze how the proposed rules will affect the value of the Internet—understood as both the physical communications network as well as the content that travels over that infrastructure. With this framework, they examine the tradeoffs inherent in Internet policy and point the direction toward rules that will facilitate the growth of the Internet and give private companies the correct incentives to continue investing in this significantly valuable good.

The explosive growth of the Internet has fundamentally redrawn our economic and social lives. The primary goals of the FCC’s Internet policy should be to ensure that the growth of the Internet continues, and to protect the innovation, creativity, and sharing that the Internet empowers. By carefully examining the likely effects of the FCC’s proposed rules, Chettiar and Holladay have made an extremely valuable contribution to the debate and to the future development of this vital national resource.

Richard L. Revesz
Faculty Director
Dean, NYU Law

Michael A. Livermore
Executive Director

Executive Summary

By giving players the best incentives for optimal investment, net neutrality encourages a cycle that breeds more content, which in turn breeds more users.

The Federal Communications Commission has proposed rules that would keep the doors to the Internet open. The proposed rules would make net neutrality the law of the land, ensuring that the Internet remains free and open to content providers.

This Report analyzes federal net neutrality policy from an economic perspective in order to understand the fundamental tradeoffs: How do we maximize the value of the Internet? Who wins, and who loses?

While there may be other considerations for policymakers, economic criteria can indicate which policy maximizes net benefits for society, and, thus, are clearly important.

There are five core findings of this Report that should influence the debate over net neutrality.

Internet Market Failure

The Internet—understood both as the physical infrastructure as well as the content and information moving along that infrastructure—produces billions of dollars of free value for the American public: Information is shared, reused, and reconfigured without fees or penalties. Websites are not compensated when their content is repurposed or passed on—that means fewer subscriptions to paid services, fewer direct page views, and a loss of advertising dollars. This economic dynamic has been taken for granted as the Internet grew around the idea that information resides in the public domain—free to be emailed, Tweeted, blogged, and discussed.

Smart Policy Can Help

As a result of this dynamic, the Internet is more useful to everyone on it, but Internet Service Providers (ISPs) and content providers are at a disadvantage since they are not compensated for all the information they disseminate. This leads to systematic underinvestment in the Internet: if that income could be accessed, it would encourage investment in infrastructure and content. Government policy can increase investment in both Internet content and infrastructure to overcome this market failure.

Transferring Wealth Through Price Discrimination

Without net neutrality rules, new technologies could lead to pricing practices that transfer wealth from content providers to ISPs, a form of price discrimination that would reduce the return on investment for Internet content—meaning website owners, bloggers, newspapers, and businesses would have less incentive to expand their sites and applications.

Efficiently Supporting Infrastructure

Additional investment in broadband infrastructure would also increase the value of the Internet—making it faster and accessible in more places. But charging content providers for access to ISP customers is an extremely inefficient economic tool to do that, primarily because most additional revenue generated for ISPs is likely to be transferred to their shareholders rather than invested in expanding broadband lines.

It is relatively easy to directly support infrastructure development, but hard to provide direct support for content. Targeted government support for ISPs to expand access where needed, along with net neutrality rules to protect content providers, are the best combination of policies for overcoming the market failure of underinvestment in the Internet.

Problems With Prioritization

Without net neutrality rules preventing priority pricing techniques, there could be changes in the way content appears online. If ISPs create “priority” or “fast lane” access to content providers at a fee, users could experience uneven access to websites and applications. While some content providers may benefit from this architecture, many type of websites will be especially harmed. Ultimately, prioritization could reduce incentives for content creators, potentially lowering the overall value of the Internet for all users.

By giving players the best incentives for optimal investment, net neutrality encourages a cycle that breeds more content, which in turn breeds more users. A combination of policies that protect content providers and judiciously deploy government resources to augment private investment in physical infrastructure is the right mix to ensure that the Internet continues to grow and flourish, generating massive benefits for the American public.

Introduction

The Internet Policy Landscape

Network neutrality (commonly referred to as “net neutrality”) as applied to the Internet allows end-users (everyday users of the Internet) the “freedom to access the content, services, applications, and devices of their choice.”¹ This seemingly simple concept has many different applications in practice. Perhaps the most important is the principle that Internet Service Providers (ISPs, the companies that own the Internet network infrastructure) must charge the same price for all content providers (companies or individuals that post websites) who wish to reach the ISPs’ subscribers.

Currently, ISPs (such as Time Warner and Verizon) operate under a *de facto* net neutrality regime: they do not charge content providers (such as Yahoo and Wikipedia) for access to their subscribers. Instead, content providers only pay fees to a single ISP to upload information to the web.² But there is no federal law preventing ISPs from charging content providers different prices for access to their subscribers. The Federal Communications Commission (FCC) has favored nondiscrimination in its policy statements, and it is now attempting to codify net neutrality principles to bring clarity to the bounds of ISP action. Under a legal net neutrality regime, ISPs would be prevented from price discriminating against content providers.

If ISPs could price discriminate, they would be able to charge content providers different prices to reach their subscribers. For example, a company like Time Warner would be able to charge a content provider like Yahoo a different price than Wikipedia to reach Time Warner’s subscribers. Or, if Wikipedia could not reach an agreement with Time Warner, then Time Warner’s subscribers would be unable to access the site. This prospect threatens the end-to-end architecture that has made the Internet so valuable: currently, all users can access all content on the Internet. The alternative to an end-to-end architecture would be a proliferation of fiefdoms of ISP-specific networks in which some sites might be available to all Internet subscribers while others might only be available to subscribers of a particular ISP. This result is referred to as network fragmentation.

The Current Process

The FCC's rulemaking proceeding initiated in October 2009 will affect which of these regimes should govern the Internet going forward. In its proposed regulation, the FCC has chosen to enshrine the current *de facto* net neutrality regime into the law. This rulemaking proceeding is part of the agency's ongoing effort to establish Internet policy.

Established in 1934, the FCC is an independent agency of the federal government.³ It is authorized to regulate all use—by anyone other than the federal government—of the radio spectrum (including radio and television broadcasting), all interstate telecommunications (including wire, satellite, cable, and Internet), and all international communications originating or terminating in the United States.⁴

“Net neutrality” as applied to the Internet mandates that Internet Service Providers (ISPs, companies who own the Internet wiring) must charge the same price for all content providers (companies or individuals that post websites) who wish to reach the ISPs' broadband subscribers.

Net neutrality prevents ISPs from favoring certain content providers over others.

The FCC typically exercises its broad jurisdiction over communications through case-by-case adjudications, rather than through rulemakings or regulations.⁵ Historically, Internet policy has been no exception; the FCC has preferred to adjudicate disputes regarding federal Internet policy on a case-by-case basis.⁶

In 2005, the FCC adopted a Policy Statement clarifying the scope of its authority to regulate the Internet, and expressing a desire to promote net neutrality on the Internet.⁷ That Statement outlined four guiding principles aimed at preserving the openness of the Internet, fostering creation and innovation regarding broadband information and services, and preserving competition and innovation to benefit users:⁸

- users are entitled to access the lawful Internet content of their choice;
- users are entitled to run applications and use services of their choice, subject to the needs of law enforcement;
- users are entitled to connect their choice of legal devices that do not harm the network; and
- users are entitled to competitions among network providers, application and service providers, and content providers.

In 2007, the FCC published a Notice of Inquiry that recommended the FCC “add a new principle to [its Internet] policy statement to address incentives for anti-competitive discrimination and to ensure the continued vibrancy of the Internet.”⁹ While that rulemaking was pending, the FCC formally committed to its net neutrality policy in its well-known “Comcast Order,” issued in August 2008. There, the FCC held that Comcast violated the FCC's Policy Statement and illegally blocked access to lawful online content by delaying subscribers' downloads and blocking their uploads. The Commission found that Comcast failed to exercise “reasonable network management” when it “selectively” targeted and interfered with BitTorrent, a peer-to-peer (“P2P”) application (a website with software allowing users to download information uploaded by other users without central coordination).¹⁰ Specifically, Comcast monitored its connections to identify P2P connections, and if it determined that “too many” customers were uploading P2P, it would drop the P2P connections.¹¹

The FCC confirmed its stance to enforce the principles that users should be able to access any content and any application, and warned that, “[i]f in the future evidence arises that any company is willfully blocking or degrading Internet content, affected parties may file a complaint with the Commission.”¹² Ultimately, the Comcast Order held that an ISP may not degrade or block legal content in the name of “network management” unless the ISP shows that its network management practice is reasonable. The FCC found Comcast’s blocking of BitTorrent to be unreasonable, as it had arbitrarily picked an application and blocked its subscribers’ access to it. The FCC held Comcast in violation of the FCC’s Internet Policy Statement because it prevented users from “running applications of their choice.”¹³ Comcast has appealed this decision to the U.S. Court of Appeals for the District of Columbia, arguing that the FCC has no authority to enforce a Policy Statement.¹⁴

The FCC’s proposed regulation, called “Preserving the Open Internet,” will affect how the Internet will be governed going forward. It codifies four existing principles of net neutrality and adds two new principles of nondiscrimination.

If finalized, the current *de facto* net neutrality system will be enshrined into the law.

On October 22, 2009, the FCC formally proposed a net neutrality regulation, entitled “Preserving the Open Internet.” that builds on the FCC’s earlier adjudication-based policymaking for the Internet.¹⁵ The FCC explains its choice to promulgate rules in this area: the Internet market has developed to the point where codified rules can “establish clear requirements” for parties while also providing the FCC with the “flexibility to consider particular circumstances case by case.”¹⁶ The rules are codified at a “relatively general level and leave more detailed rulings to the adjudications of particular cases.”¹⁷ To enforce these rules, the FCC can initiate an enforcement action on its own, or in response to a complaint filed by a third party.¹⁸

The new regulation applies to all “providers of broadband Internet access,” defined fairly broadly to include all ISPs providing Internet access service (other than via dialup.)¹⁹ to end-user subscribers. The FCC seeks comment on whether and how to apply this regulation to “managed or specialized services” (such as subscription voice and video services, applications for telemedicine, smart grid, eLearning, or cable television) and to wireless services (including mobile wireless).²⁰

The regulation first codifies the four net neutrality principles outlined in the FCC’s 2005 Policy Statement. Notably, the principles are rewritten as rules that impose obligations on broadband ISPs rather than as entitlements of users:²¹

1. *Subject to reasonable network management, a provider of broadband Internet access may not prevent any of its users from sending or receiving the lawful content of the user’s choice over the Internet.* This “content choice rule” ensures that users can both send and receive the content of their choice, expressing their views.
2. *Subject to reasonable network management, a provider of broadband Internet access may not prevent any of its users from running the lawful applications or using the lawful services of the user’s choice.* This “applications and services choice rule” ensures that users can run applications and use services of their choice.
3. *Subject to reasonable network management, a provider of broadband Internet access may not prevent any of its users from connecting to and using on its network the user’s choice of lawful devices that do no harm the network.* This “devices choice rule” allows users to connect to the Internet by using their choice of legal devices that do not harm the network.

4. *Subject to reasonable network management, a provider of broadband Internet access may not deprive any of its users of the user's entitlement to competition among network providers, application providers, services providers, and content providers.* This “provider choice rule” protects competition among network providers, application and service providers, and content providers. The rule does not define “application, content, or service provider,” as any user of the Internet can be such a provider.

The rulemaking also proposes two additional principles:

5. *Subject to reasonable network management, a provider of broadband Internet access must treat lawful content, application, and services in a nondiscriminatory manner.*²² This “nondiscrimination rule” intends to prohibit the ability of network operators to discriminate in price or service quality among different types of traffic or different providers or users, unless that provider is able to prove such discrimination is necessary for “reasonable” network management.
6. *Subject to reasonable network management, a provider of broadband Internet access service must disclose such information concerning network management and other practices as is reasonably required for users and content, application, and service providers to enjoy the protections specified in this part.*²³ This “transparency rule” allows all of the public (especially content, application, and service providers, users, the FCC, and other parties) to review and understand the policies of ISPs. Those parties can then bring enforcement actions if it appears the policies are in violation of any of the FCC’s rules.

The nondiscrimination rule, which is similar to what the FCC recommended in its 2007 Notice of Inquiry, is most relevant to the net neutrality debate. This rule would prohibit a broadband ISP from discriminating against, or in favor of, any content, application, or service, subject to “reasonable network management” (explained below).²⁴ The FCC notes that the nondiscrimination rule prevents an ISP from “charg[ing] a content, application, or service provider for *enhanced or prioritized access* to the subscribers of the broadband ISP.”²⁵ However, the rule “would not prevent a broadband Internet access service provider from charging *subscribers* different prices for different services.”²⁶

The nondiscrimination rule, like all the proposed net neutrality rules, would allow ISPs to engage in “reasonable network management,” defined as “reasonable practices” by an ISP in order to: (i) combat “effects of congestion” or “address quality-of-service concerns”; (ii) address “unwanted” or “harmful” traffic; (iii) prevent transfer of unlawful content; or (iv) prevent unlawful transfer of content.²⁷ In an explanation of the first part of this definition, the FCC explains that what constitutes “congestion” and “quality-of-service concerns” may vary depending on the situation, and the agency gives several examples of what may qualify: temporarily limiting bandwidth usage of individual users whose use is affecting other area users, or charging users based on actual bandwidth usage rather than by a flat fee.²⁸ The first example indicates that the FCC may allow ISPs to treat individual users differently, as long as not on the basis of content and the treatment improves other users’ connections.

Regarding the second part of the definition of “reasonable network management,” the FCC gives three examples. The first two, blocking spam and malware (malicious traffic), seem to align more closely with the “harmful traffic prong,” while the last example, “any traffic that a particular user has requested be blocked,” seems to speak to a type of information unwanted by users.²⁹ As to the third part of the definition, the FCC clarifies that it covers not only illegal distribution but also unlawful distribution of copyrighted material.³⁰

The FCC notes that the fourth part of the definition is meant to serve as a “catch-all” prong; specifically, it intends to give ISPs some flexibility to use “reasonable network management” to respond to unanticipated changes in technology or broadband use.³¹ Finally, the regulation affirms it does not modify or conflict with ISPs’ current legal obligations—for example, providers must still allow for the needs of law enforcement, public safety, and national security.³²

This is the net neutrality policy now on the table.

An Economic Framework

From an economic standpoint, the goal of federal government Internet policy should be to maximize the net present value of the Internet. The Internet can be thought of as a collection of productive economic assets: the combination of the intellectual property, physical capital, and human labor that has been and is devoted to building and maintaining the physical structure and content that comprise the Internet. Maximizing the net present value—that is, the present value of all net future flows of revenues (or, more broadly, utility) derived from the Internet—ensures that these resources are used in the most productive manner possible. Ensuring an optimal level of investment and optimal allocation of those investments will maximize the net present value of the Internet.

“Price discrimination” is an economic concept that occurs when a seller offers the same good to different consumers at different prices based on the amount a consumer is willing to pay. Currently, websites only pay to upload information, not when their information is downloaded.

Price discrimination would allow ISPs to charge content providers different prices when the ISPs’ subscribers download information from content providers.

Both proponents and opponents of net neutrality employ economic arguments to advance their positions. Unfortunately, the necessary data to precisely estimate the “value” of the Internet, or the costs and benefits of net neutrality rules is not available.³³ However, economic theory can be used to estimate the impact of policy changes on the value of the Internet. On balance, because they best align the incentives of private parties with optimal investment in the Internet, net neutrality rules are likely to be economically justified.

Part One of this Report explains how the Internet—understood both as the physical infrastructure as well as the content and information moving along that infrastructure—produces billions of dollars of free value for the American public. These “positive externalities” lead to systematic underinvestment in the Internet and may justify government intervention to overcome this market failure and incentivize the optimal amount of investment in the Internet. Part Two discusses how the core of the net neutrality debate is about a tradeoff of wealth: eliminating net neutrality would allow ISPs to implement new technologies to institute pricing practices that would transfer wealth from content providers to ISPs.

Part Three analyzes how this wealth transfer affects incentives to invest in different parts of the Internet: giving more wealth to ISPs would allow them to use some of that revenue to make additional investments in the Internet infrastructure; but that would come at the expense of content providers, who would have less incentive to invest in Internet content. Further, most of the additional revenue generated by ISPs’ widespread use of price discrimination or traffic prioritization would compensate for ISPs’ past investments, and ISPs would probably transfer that revenue to their shareholders rather than use it to investment in new broadband infrastructure. Transferring wealth to ISPs in no way guarantees that ISPs would invest that wealth in Internet infrastructure. It is also likely ISPs would not be able to perfectly price discriminate, leading to

inefficiencies that would decrease the value of the Internet. Part Four explains how government can correct the positive externalities in the Internet market by supporting different parts of the Internet. It also discusses how difficult it would be for government to subsidize Internet content providers compared to how easy it would be for government to subsidize Internet infrastructure.

Part Five turns to the issue of “prioritization” pricing schemes, such as “fast lane” architecture. Like price discrimination, these

prioritization schemes would result in underinvestment in many forms of Internet content. While some content providers may see net benefits from this type of architecture, forms of content that are not easily monetized would be especially harmed by the creation of a fee-based fast lane access to ISP customers.

Overall, net neutrality maximizes the value of the Internet. Opponents of net neutrality do have legitimate goals, but those goals can be accomplished more efficiently through other government policies that do not harm the functioning of the network on which so much of our economy relies. If policymakers’ overall objective is to protect the current structure of the Internet while increasing broadband penetration and speed, the most beneficial policy proposal would be to enshrine net neutrality in the law to protect content providers while creating separate government programs to directly support expansion of the physical Internet infrastructure.

From an economic standpoint, the goal of FCC policy should be to maximize the net present value of the Internet: the combined value of the Internet infrastructure and the content that flows to end-users.

This Report concludes that, all things considered, enshrining net neutrality as law would likely lead to a higher value of the Internet.

Part One

Positive Externalities: The Hidden Value of the Internet

In a perfect world, markets would always be efficient. In the real world, however, there are many types of “market failures” that prevent the allocation of scarce social resources to where they would create the highest value. In such cases, government must sometimes step in to recalibrate and level the playing field, generating economic benefits for society.

This Part explains the structure of the Internet and how that structure generates “positive externalities,” a particular kind of market failure. These positive externalities create a market in which producers are not fully compensated for the benefits their goods create. As a consequence, producers will end up underproviding the good—production therefore falls below what would be socially optimal.

A. The Internet Structure

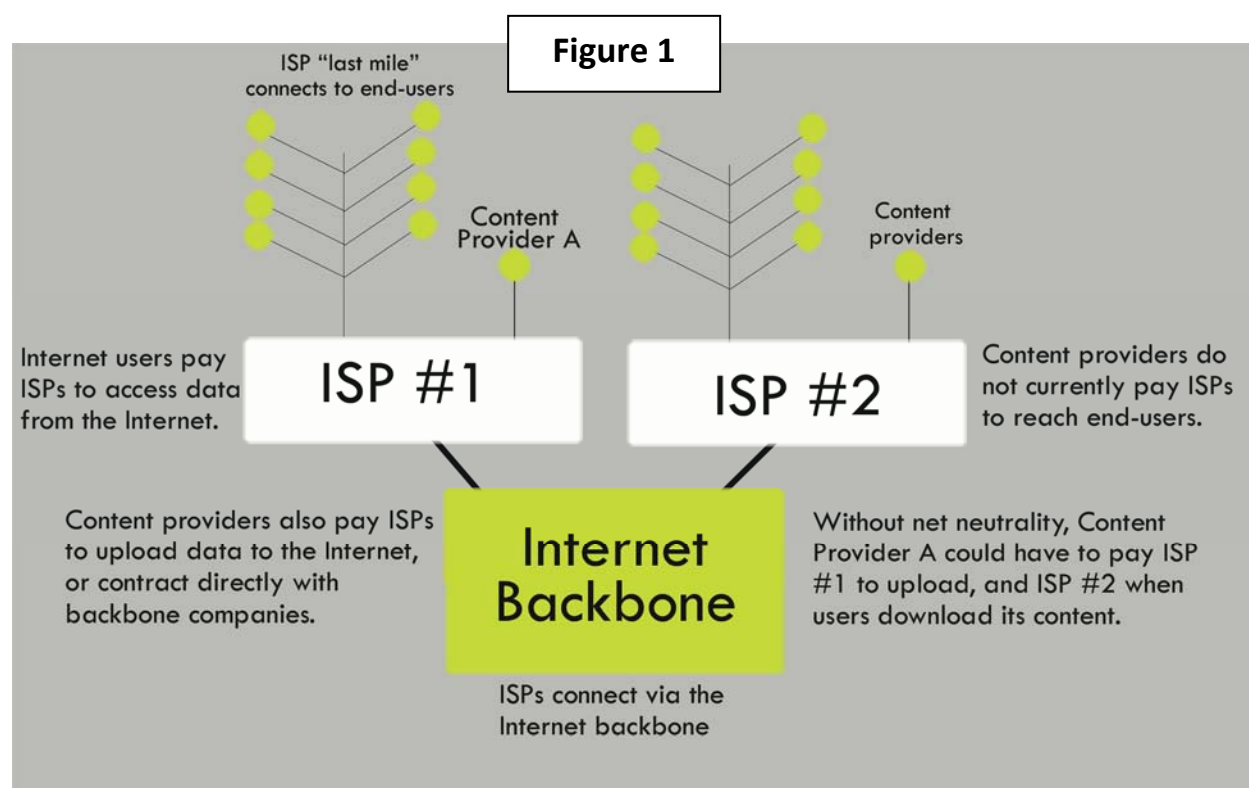
The Internet can be thought of as two complementary goods: infrastructure and content. Internet infrastructure includes the broadband lines that run across the country (and the world), connecting users and content providers to ISPs, ISPs to other ISPs, ISPs to backbone companies, and backbone companies to each other. It also includes the satellite and cellular towers that provide wireless Internet; the ISPs’ “points of presence” (the end of the ISPs’ networks); the millions of computers, cell phones, televisions, servers, and other pieces of hardware connected to the network; and the telephone lines over which some users connect to the Internet. Internet content, on the other hand, includes all the information, websites, blogs, videos, applications, software, et cetera found on the worldwide web.

Internet subscribers obtain information and experiences from websites, applications, and other content types. This information (Internet content) provides some value to the Internet user. Because of the nature of information, it can often be shared for free. For example, a user can read an entry on broadband policy on a content provider like Wikipedia’s website and then share that

information with someone else. This secondary sharing is an important piece of the Internet's value to society. Internet subscribers clearly receive a great deal of benefit from content and applications that run over the network, but people also enjoy informational benefits from others' use of the Internet. In addition, the more people accessing, using, and adding information to the network, the more valuable a connection to the network becomes for everyone.

Network access (Internet infrastructure) is also an important input for a great variety of industries, so high quality Internet access is critically important to the economy as a whole.³⁴ In addition to a large consumer market for Internet access, the largest and fastest growing sectors of the United States economy use Internet access as an input into their production processes and to deliver information to their users. Service sectors—including educational services; professional and business services; and finance, insurance, and real estate industries—use Internet access to provide information in a timely manner to both employees and customers. The ability to communicate quickly and flexibly over the Internet is at the heart of the U.S. economy's structure.

Figure 1 provides a simple representation of the structure of the Internet:



On the left side of the diagram are content providers who upload their applications and websites onto the Internet via a transit provider. Content providers pay a fee to the transit provider—often an Internet Service Provider (ISP), but it could be any of a variety of types of companies that sell access to the Internet. This is typically the only fee that content providers pay to access the Internet and Internet subscribers. ISPs, which provide services to both content providers and end-users, connect their private networks to the Internet backbone in the center of the figure. The connections between an ISP's Internet access point and its subscribers' computers constitute a privately owned subnetwork: the ISP owns and operates this network, which is often referred to as the "last mile" of the Internet, connecting the network to individuals.

The last mile is at the heart of the net neutrality debate. The cost of building a last mile network is extremely high and is often borne entirely by the ISP that constructs the network. Building this type of network requires ISPs to build physical or wireless connections between their Internet access point and each subscriber's household or business. This last mile network is the ISP's most valuable asset.³⁵ The ISP's router acts as a "gate" from the larger Internet network to its private last mile network.

ISPs connect from their networks to Network Access Points ("NAPs"), which connect different ISPs' networks to each other. The NAPs around the world are connected to each other through the backbone infrastructure, creating the Internet. Internet backbone companies own this infrastructure, which moves data between different NAPs. Common examples of Internet backbone companies are Level 3, Qwest, Tata Telecom and Reliance Communication—they deal in the business-to-business market and typically do not market themselves to the public. While not all household names, they provide the crucial service of connecting ISPs across the country. Together, these players make up "the Internet."³⁶

B. Internet Market Inefficiency

Critics of net neutrality sometimes argue that there is no role for government in the Internet market. In general, it is true that market mechanisms tend to be the best tool for determining the optimal level of investment, the best prices, and the correct quantity of available goods. Opponents of net neutrality argue that, because net neutrality involves government intervention in the market, it "could substantially reduce investment incentives, distort innovation, and ultimately harm [users]."³⁷ They argue that creating rules that govern how ISPs manage their network could reduce the functionality of the Internet. ISPs would like the ability to manage traffic on their networks in any way they see fit, and they fear that net neutrality regulations would tie their hands from doing so.

Opponents also claim that limiting ISPs' ability to charge content providers and application programmers for access to the ISPs' subscribers decreases the incentives to invest in the network. If ISPs are limited in their ability to earn revenue from their investment in network infrastructure, then they would have less incentive to invest. In turn, ISPs' underinvestment in building network infrastructure could starve the Internet for the bandwidth necessary to provide the new voice, video, and other technologies currently being deployed on the Internet.

Some also fear that the government cannot respond to the quickly changing technological developments and reshape regulation to new Internet developments. The government has a decidedly mixed history in imposing price regulation, and the Internet is a crucial component for many industries in the United States economy. For this reason they argue that regulating ISPs and the Internet more generally is simply too risky.

However, all these arguments fail to grapple with the fundamental fact that the Internet does not function like a theoretical market: due to its network structure and the ability of users to transfer knowledge for free to nonusers, the Internet creates pervasive positive externalities that

An "efficient market" maximizes total surplus—together buyers and sellers reap the maximum amount of surplus from the market. By maximizing total surplus, the market achieves the greatest possible benefit to society in the aggregate.

Efficiency does not determine the distribution of surplus.

systematically reduce incentives to invest. In these circumstances, well-designed government policy may be warranted to correct this market failure.

Market Efficiency Theory


In theory, markets should reach the point of efficiency on their own. A market generates economic benefits when a consumer and supplier voluntarily enter into a transaction. A “reservation price” is the highest price a consumer is willing to pay, and the lowest a supplier is willing to accept. The market price is, by definition, lower than the consumer’s reservation prices and higher than the supplier’s reservation price.

The difference between the reservation prices of the seller and buyer and the actual price is called an “economic surplus.” A consumer willing to pay \$10 for a good, who actually pays \$5, sees an economic surplus of \$5. Total economic surplus represents the benefit to both the consumer (consumer surplus) and producer (producer surplus) of entering into this transaction.


To be efficient, a market must maximize total surplus. Notably, efficiency does not determine the distribution of surplus; there are efficient markets, for example, in which consumers receive the entire economic surplus. However, as explained in Part Two, the debate about net neutrality revolves around how the benefits of Internet access should be allocated between consumers (here, content providers who consume access to Internet end-users) and producers (ISPs that generate access to Internet end-users).

Access to information on the Internet generates large benefits for both end-users and content providers. Each time a consumer uses the Internet to explore content or use an application, there is a potential for economic surplus to accrue. For example, an Internet user may find out about a sale on a store’s website, allowing the user to save money on a purchase and the company to earn an extra customer. Both the user and store receive a benefit from this transaction.

The traditional supply-and-demand graph illustrates how markets generate and distribute economic surplus and reach optimal levels of production. Figure 2 compares demand and supply. The intersection of the supply and demand curves determines both the market price and the quantity of production of the good. In the absence of market failure, this price will generate the largest possible economic benefits from buying and selling this product. The blue triangle represents consumer surplus—the benefits to consumers from participating in the market. The red



triangle represents producer surplus—the benefits that accrue for producers from participating in the market. At the optimal price and quantity, the combined area of the two triangles is maximized, creating the maximum combined surplus.



The Internet access market should generate surpluses for both Internet users and Internet Service Providers in this same way. If the market is functioning properly, the price for Internet access will maximize the total surplus and distribute that surplus between producers and consumers. If the market price rises above the efficient level, some customers will be priced out of the market, eliminating their consumer surplus.

Figure 3 shows what happens if prices are out of equilibrium. In this example, some users would no longer purchase Internet access, and consumer surplus would shrink to the area of the blue triangle. The grey

triangle represents the lost surplus under the new pricing plan, and the wealth in this area is lost to both Internet users and ISPs. It is this loss that economists seek to avoid by encouraging the use of markets, instead of governments, to set prices.

The impact of the increase in price on producer surplus is unclear. Producers lose part of the area of the grey triangle, but gain an increase in the size of red polygon from the increase in price. This surplus increase may more than offset the decrease from a reduced number of users. In that case, ISPs can even benefit from a price increase that reduces the overall surplus generated by the market.


Market Failures & Inefficiencies

The Internet market is not efficient for two primary reasons: informational externalities and network externalities. “Externalities” occur when the purchase of market goods leads to benefits or costs that are not captured by the original buyer or seller. In other words, the transaction produces effects that are external to the market. This is a type of market failure that leads to inefficient market outcomes.

Pollution from a smokestack is a famous example of an externality. When power plants burn coal to generate electricity, they also generate smoke laden with pollutants. That smoke falls far away from the power plant and causes health problems, reduced visibility, and acid rain. The individuals harmed by these negative outcomes are not compensated by the power plant that profits from the sale of electricity or by the buyers who benefit from their use of electricity. The costs of those

damages are external to the market, and so the power generator does not take those damages into account when making their business decisions. This is an example of a negative externality that imposes costs not captured by the market. Markets tend to overprovide goods with negative externalities because the external damages are not considered by the producer when making its output and price decisions. If the electricity generator were considering all the damages caused by its emissions, it would produce less electricity and correspondingly less pollution.

Positive externalities are the opposite—in these cases, there is some positive effect from the good that is not captured by market participants, either consumers or producers. Markets generally tend to underprovide goods that create positive externalities. Figure 4 illustrates the impact of a positive externality on a supply-and-demand framework. The external benefit does not affect the supply curve, but it splits the demand curve into two.



The external benefit does not affect the supply curve, but it splits the demand curve into two.

One demand curve represents the private benefits, and the other demand curve includes both the private benefits and the external benefits. The distance between the two demand curves marks the size of the externality. The market equilibrium is where the supply curve crosses the private demand curve; but the optimal outcome from society's perspective occurs where the supply curve crosses the social demand curve. The market provides too low a quantity and charges too low a price relative to the social optimum. In this way, externalities result in an inefficient market, creating the opportunity for government intervention to

increase total surplus.

Network Externalities

The Internet market exhibits a network externality—also called the “network effect.” This effect occurs when the value of a good or service increases as other users purchase the same good or service.³⁸ A telephone is an example of a good that exhibits network effects: a single person who purchases a phone does not reap benefits unless other users buy telephones and use them to complete calls.

The network effect is conceptually similar to “economies of scale,” in which the costs of producing a good decrease with output. Typically, marginal costs increase with production—the more goods produced, the greater the additional cost to create each unit. Goods with economies of scale exhibit


the opposite trend: as more units are produced, marginal costs go down. Industries that have high overhead costs and low unit costs of production tend to exhibit economies of scale. For example, provisioning of tap water and automobile production are both industries that exhibit strong economies of scale. Economies of scale, however, occur on the production side of the market; the network effect acts on the demand side of the market.

Network effects cause the value of a network to increase with the number of connections in the network. That means the Internet becomes more valuable to each user as additional users are added. This phenomenon occurs in part because Internet users are also content creators. Few people can program Java apps, but many Internet users comment on news stories, send emails to listservs, and post items to Craigslist and eBay. The structure of the Internet makes it fairly easy for each consumer to generate content that other users will find valuable.

If a potential user does not connect to the Internet because she does not find enough valuable information on it, the value of connecting to the Internet is reduced for everyone. Individual websites, such as eBay and Craigslist, also benefit from the network effect—as their number of users increase, their websites become more valuable to users, creating a positive feedback loop. If eBay had a single user, it could not function. The more users who join the site as buyers, the more sellers are attracted to the site to place their goods. Similarly, as more sellers join the site, potential buyers know they are more likely to find the items they seek, and the site becomes more attractive from a buyer's perspective. As more users of either type join the site, it becomes a better resource for all users.

These powerful feedback loops characterize networks and make competing against existing networks particularly difficult. For example, America Online famously built a closed network for their America Online Instant Messaging service (AIM), which allowed users to communicate in real time only with other AOL customers.³⁹ By closing its network, AOL reduced the value of AIM to their own users and also decreased the value of competitors' chat networks (such as those of Yahoo and Microsoft).

These effects can also make networks extremely valuable and give existing networks a built-in advantage over competitors. An upstart company attempting to develop an Internet chat system or online marketplace must market itself widely to gain the necessary number of users to make its network valuable to users. Existing companies can rely on the value of their network to keep users satisfied and encourage new users to join. As networks expand, this feedback system reinforces



growth. However, as networks shrink, the same feedback system works in reverse.

As shown by Figure 5, the network effect causes an unusual demand curve. Demand curves indicate how many buyers are willing to purchase a good at various prices. For a standard good, the demand curve slopes downward; this is because for most goods, there are a small number of people willing to pay a large amount to purchase the good and more people willing to pay successively smaller prices.

The network effect demand curve starts somewhat lower than the standard demand curve because without other people on the network the willingness of a consumer to purchase the good is lower. Instead of sloping downward, the network effect demand curve slopes upward. This is because as more people join the network, users' "willingness-to-pay" for the good increases. At a certain point the standard demand effect kicks in and additional end-users are not willing to pay as much for access to the network. At that point, the network effect demand curve also begins to slope downward. The size of the hump in the demand curve illustrates the strength of the network effect. The higher the hump, the stronger the network externality—and the more a government must consider network effects when considering policy in that market.

The Internet market is not efficient. An "externality" is a type of market failure which occurs when the purchase of goods leads to benefits or costs that are not captured by the original buyer or seller.

On the Internet, the ability to reuse, share, and expand on information is a positive externality to users.

Because of the Internet's network effect—which causes a market failure just like any externality—the Internet market will tend to underprovide content and infrastructure relative to the socially optimal level. Given the strong relationship between the Internet's size and its value, any policy action that reduces the number of Internet subscribers or the quantity of content on the Internet is likely to be damaging. FCC policy should take into account the impacts of the network effect when estimating the impact of any policy on the value of the Internet.

Information Externalities

The information on the Internet exhibits another positive externality because it can be spread for free. Many types of information are "public goods," meaning that they are both nonrival and non-exclusive. To be nonrival, an additional user of a good must not diminish the good's value. A non-exclusive good is one that is difficult to prevent others from accessing. Clean air is a classic nonrival, non-exclusive public good, as everyone can enjoy clean air simultaneously without reducing its value. In the case of information, complex intellectual property rights are sometimes needed to ensure exclusivity, but there are many types of information for which these property rights do not exist.

Information first gathered online by Internet users can be reconfigured and shared with others, either on the Internet or through other forms of communication. When information is widely shared, the content provider who originally provided the information and the ISP who made it possible to access the information will not necessarily be compensated for any additional information exchange. But those exchanges, as information is passed around, provide value to the individuals who received the information. This is a well-known phenomenon known as the "information externality."⁴⁰ Because information providers are not fully compensated for the value of the information they provide, the free market tends to undersupply information for the same reasons described above when any type of positive externality occurs.

The Effects of Externalities

Some observers of the Internet market rely on textbook economic theory that government involvement in the market leads to inefficiencies.⁴¹ This argument is usually correct. As discussed above, in the absence of market failure, the equilibrium market price will generate the optimal amount of production and an overall efficient outcome. A well-functioning market optimally sets price and quantity of the good, and ensures that the sum of the producer and consumer surplus is as large as possible. This economic theory generally argues that price regulation reduces the efficiency of a market by constraining the behavior of sellers.⁴² Any government action in such a market that changes the price or quantity of production will lead to inefficiency.⁴³

In the case of a market that exhibits externalities, however, that argument no longer holds. The presence of these two positive externalities ensures that the Internet market, left to its own devices, will produce inefficient outcomes and underprovide both infrastructure and content. If the aim of government policy is to maximize the value of the Internet, price regulation may be justified to address the network and information externalities. The necessity of fully compensating content providers for the informational benefits they provide to society and correcting the network effect is well known, but there are few easy solutions. When the true value of this information is unknown, the true value to society of investments in content and the infrastructure necessary to access that content are hidden. The private benefits of investment in the Internet are lower than the social benefit because the network and information externalities make it impossible for Internet companies to receive compensation for all the benefits they provide. As long as these externalities are uncorrected, Internet content and infrastructure will continue to be underprovided.

The Internet does not function like a theoretical market. Due to its network structure and the ability of users to transfer knowledge for free to nonusers, the Internet creates pervasive positive externalities that systematically reduce incentives to invest.

In these circumstances, well-designed government policy may be warranted.

Economic literature suggests that markets with network effects that underprovide goods can be corrected through government subsidization.⁴⁴ The government may reduce the costs of entering a network for early adopters to encourage entry into the network. In effect, because the early adopters produce positive externalities, the government compensates them for participating in the network, either directly through transfer payments or through some other mechanism. This can overcome the “chicken and egg” problem facing goods that exhibit a network effect: they cannot gain users unless they become more valuable, but they cannot become more valuable unless they gain more users. Government assistance for early adoption can help networks reach a sufficient size to attract new subscribers independently.

Supporters of price discrimination argue that because Internet service provision is a two-sided market, ISPs should be able to charge both sides of that market: charge users for access to the Internet and charge content providers for access to users.⁴⁵ However, in two-sided markets with network effects and other positive externalities, it may be the case that *both* sides of the market should be *subsidized* rather than charged.

If FCC policy can respond to the externalities in the Internet market, it will increase the value of the Internet to society as a whole. A beneficial government policy to correct the under-provision of such vital goods should bring the market back to the correct equilibrium by incentivizing market players to provide both more content and more infrastructure. In the following discussions, this

Report describes a proposed policy that can do just that—through a combination of direct subsidies for Internet infrastructure, and net neutrality to protect the surplus for content providers. Taken together, these measures correct for at least some of the positive externalities associated with the Internet by raising the overall level of investment compared to an unregulated market.

Part Two

The Net Neutrality Wealth Tradeoff: Content Providers and Internet Service Providers

The government's choice of whether to adopt a net neutrality framework will have serious consequences for how the economic surplus from the Internet market is distributed between content providers and ISPs. In our current *de facto* system of net neutrality, content providers enjoy large portions of the surplus, and thus, receive increased revenue and rates of return on investment. If we were to change over to a system of price discrimination, revenue currently in the hands of content providers would be transferred into the hands of ISPs. Currently, ISPs do not charge content providers to reach the ISPs' subscribers. Under a system of price discrimination, ISPs would begin charging content providers to reach their subscribers and could charge different content providers different amounts—either for the same service or for priority service (namely, a faster connection to subscribers). Because content providers would have to pay ISPs for this service, part of the revenue generated by content providers would be transferred to ISPs. This wealth transfer will put more money in the hands of ISPs, increasing the rate of return on infrastructure investment, but lowering the rate of return on content investment.

The choice of government policy will dictate who keeps what portion of the revenues generated in the Internet market. Who receives this surplus does not necessarily affect the efficiency of the market. However, in the context of the Internet, this distributional question—in combination with both the positive externalities produced by both content providers and ISPs (explained in Part One) and the relative difficulty of directly subsidizing content providers versus ISPs (explored in Part Four)—can have important efficiency implications.

This Part will first discuss the current pricing structure that exists under the *de facto* net neutrality regime and will explain how new technologies have allowed new pricing models. It will also describe in more detail the concept of price discrimination and how it can be used to effectuate a wealth transfer from content providers to ISPs. As explored in Part Three, who “wins” this tug-of-war over economic surplus will have important consequences for incentives to invest in different parts of the Internet.

A. Current Internet Pricing

There are many different forms of pricing involved in the Internet market: users pay ISPs; content providers pay ISPs or backbone companies; ISPs pay each other; ISPs pay backbone companies; and backbone companies pay each other. The net neutrality debate most directly affects the prices that ISPs charge content providers to use the ISPs' last mile (broadband lines that connect to the ISPs' subscribers).

Everyone, from single users to large companies like Google, must pay for access to the Internet. Currently, all subscribers (including Internet users, content providers, and businesses) pay to connect to the Internet usually using ISPs' broadband networks. Household fees are typically structured as a monthly charge for unlimited access to the Internet. Large businesses typically pay fees based on the quantity of data they upload and download from the network.⁴⁶ Internet content providers contract with ISPs or backbone companies based on the quantity of data uploaded. The market for uploading content is relatively competitive; there are multiple backbone companies and content providers can choose between ISPs and backbone companies.⁴⁷

The question of exactly how much large content-providing companies pay is difficult to answer. Google, for example, does not publish how much it pays for bandwidth.⁴⁸ As explained in Figure 1, in the current *de facto* net neutrality regime, content providers pay on the upload (they enter into agreements with ISPs to use their broadband network to upload data from their servers to the network), but they do not pay on the download to users (they do not pay ISPs that own the broadband down to everyday Internet users). ISPs could conceivably charge every content provider the same price for access to their network without falling afoul of the FCC's anti-discrimination regulation. If the ISPs were forced to charge a single price to all content providers, they would likely choose not to charge any price at all because of the large number of Internet content providers who would find the transaction costs and price of accessing an ISP's network prohibitively high.⁴⁹ Any positive price would likely reduce the amount of content on the network, frustrating Internet users and reducing ISPs' revenue from their primary customers, Internet end-users.

The current pricing scheme has raised concerns that content providers are taking advantage of ISPs' "property for free."⁵⁰ For example, when AT&T's then-Chief Executive Officer Edward Whitacre was asked about Google, Microsoft Network News, and Vonage, he stated:

What they would like to do is use my pipes [for] free, but I ain't going to let them do that because we have spent this capital and we have to have a return on it. So there's going to have to be some mechanism for these people who use these pipes to pay for the portion they're using. Why should they be allowed to use my pipes?⁵¹

In fact, however, as explained above, Google, Microsoft News, and Vonage all pay AT&T if they use its lines to connect to the Internet (if they do not pay AT&T, then they must pay whichever ISP or backbone company they choose to provide them this service). These websites do not pay to use AT&T's lines connecting from the Internet to AT&T's subscribers—the subscribers pay for their connections. As explained in the next section, ISPs now have technology that would allow them to change this pricing structure and actually charge content providers both on the way in and the way out.

In addition to the pricing between ISPs and users, there is also a pricing mechanism among ISPs. Networks pay one another for the data traffic they pass to each other.⁵² For example, if an AT&T-owned network passes 2 terabytes⁵³ of data to a Verizon-owned network and the Verizon network only passes 1.5 terabytes of data back, then Verizon will pay AT&T for the additional 0.5 terabyte of data that they passed onto the AT&T network. This same pricing technique is used between ISPs and backbone companies and among backbone companies.⁵⁴

B. New Pricing Models

The main technology that gives broadband providers the ability to price discriminate is Deep Packet Inspection (DPI) technology, which was introduced only in the last decade. Information that passes through the Internet consists of units of information called “packets.” These packets are comprised of a “header,” which contains processing information (such as the source and destination address), and a “data field,” which contains all other information.⁵⁵

Before DPI, networks processed packets using header information, as that field is instrumental in getting information to the correct destination. ISPs were thus unable to determine the source applications of the various packets that traveled over their networks. DPI changes all of this by enabling broadband providers to monitor the data field, including the source, of all incoming and outgoing packets in real time.⁵⁶

DPI has several security uses. For instance, it is particularly useful for preventing “denial of service” attacks (where hackers flood a website or computer with packets of information in an effort to slow or crash the site) because it can stop these attacks close to the source rather than at the endpoint.⁵⁷ DPI can also identify worms and other specific viruses and prevent them from infecting the network; Comcast plans to use DPI to help customers fight such problems.⁵⁸ DPI even helps in the areas of law enforcement and national security. For example, network-based DPI provides the means for telephone providers to abide by U.S. law, which requires them to be able to provide call records or call tapping to law enforcement officials.⁵⁹ Additionally, DPI allows governments to develop techniques critical to the development of cyber security or cyber espionage.⁶⁰ DPI can also be used to target advertising directly to users, which arguably should pose no privacy issues so long as the practice is disclosed and done with the permission of those affected.⁶¹ However, prior attempts at targeted advertising through DPI were abandoned due to public privacy concerns.⁶²

If an ISP can use DPI to identify and subsequently affect the transfer of individual packets in real time, then it functionally has the power to discriminate based on content, source, or both. When discriminating based on origination, an ISP could use DPI to identify the origination of a particular packet sent over its network, and could then charge different content providers different prices to send their information over its broadband lines and reach its subscribers. An ISP could also tag packets for prioritized (i.e., faster) treatment or block the packets altogether, based on what the packets contain or which application sent them.⁶³

Technologically, this means that ISPs could put packets from YouTube on a fast lane, put packets from Yahoo on the slow lane, and charge YouTube more for this fast lane treatment. Or, ISPs could charge YouTube nothing for this fast lane treatment. An ISP could also charge YouTube more for its packets than it charges Yahoo to be on the exact same lane of traffic to reach its subscribers; or it could simply refuse to carry packets from Yahoo to its users, yet carry YouTube’s packets for a higher price.⁶⁴

Alternatively, an ISP could use DPI to identify the type of packet sent over its network and discriminate based on content or application. For example, an ISP like Time Warner could charge content providers more for videos than for standard website pages to reach subscribers. It could also place videos on the fast lane, and standard web pages on the slow lane. This would mean that a company like Time Warner could charge a content provider like YouTube to place its packets in the fast lane, but not offer fast lane service to the text and images on Wikipedia. Or, a Time Warner could charge content providers more to send Peer-to-Peer (“P2P”) applications⁶⁵ over its last mile than other applications.

“Willingness-to-pay” is an economic concept that values is the maximum amount a buyer would be willing to pay, sacrifice, or exchange for a good.

The difference between a consumer’s willingness-to-pay and the actual price paid is the “consumer surplus” that is gained from the transaction.

The FCC’s nondiscrimination rule would prohibit an ISP from treating any content, application, or service in a “discriminatory” manner, subject to reasonable network management. This clearly bans pure price discrimination (charging different content providers different prices to access their subscribers). The regulation also bans ISPs from offering content providers a “take it or leave it” offer on access to their users. For example, an ISP like Verizon could not charge a website of a company like *The New York Times* a certain price for access to its subscribers by threatening to block the website from its network and therefore from its Internet subscribers.

Some ISPs are currently using Deep Packet Inspection in network management; however, none are known to be using the technology to price discriminate. In 2007, Comcast used DPI to identify P2P uploads (by determining their origination from other users) and terminate them.⁶⁶ However, the FCC subsequently found that Comcast’s actions violated its policy, and Comcast can no longer use DPI in this manner.⁶⁷ Also, Cox Communications has explored the use of DPI to identify the type of packets sent over its network and prioritize time-sensitive packets over “non-time-sensitive” ones;⁶⁸ this arguably violates nondiscrimination principles because Cox is “inserting its own value judgment [about the content] in place of a user’s judgment.”⁶⁹ Although Cox eventually decided not to implement this policy,⁷⁰ Primus Telecommunications has announced a similar strategy outside the United States.⁷¹ Some ISPs have already implemented DPI in router hardware on their network,⁷² and several companies are currently offering DPI services to United States ISPs.⁷³

Thus, with DPI comes the ability of ISPs to institute price discrimination and prioritization pricing models: ISPs can now charge different content providers different prices to use the ISP’s last mile. The availability of DPI has added immediacy to the net neutrality debate.

C. Price Discrimination Effectuates a Wealth Transfer

New technologies like Deep Packet Inspection will enable price discrimination. “Price discrimination” is an economic concept that occurs when a seller offers the same good to different users at different prices based on the amount a consumer is willing to pay for a particular good, or willingness-to-pay.

Price discrimination is not inherently bad. In its pure form, it simply represents a transfer of economic surplus and would not reduce overall surplus. For example, the processes of “bargaining” at a crafts market may lead to price discrimination as sellers attempt to guess the reservation price of buyers and refuse to sell at much below that price. If the sellers are good at their jobs, the same number of crafts is sold, but users walk away with less surplus and the sellers walk away with more surplus.

In theory, price discrimination can only exist in markets with monopolies (one seller) or oligopolies (a handful of sellers), or in markets where resale is impossible, such as medical services.⁷⁴ This is because these types of markets avoid the “arbitrage” opportunities that can be created by attempts to price discriminate. In a theoretical market, if sellers price discriminate (with perfect information, no transaction costs, and no prohibition on re-sale), individuals would engage in arbitrage. The moment the seller attempts to sell the good at a higher price, buyers who purchased the good at a lower price would resell the good to the new buyers at a price slightly lower than that offered by the seller. This arbitrage would prevent sellers from selling at the higher price, and sellers would eventually price the good at a consistent level in both markets to prevent arbitrage.⁷⁵

However, in markets where buyers cannot simply become sellers and begin selling the product, arbitrage cannot occur and price discrimination can exist. As will be explained in Part Three, local Internet service provision markets are often local oligopolies and resale of Internet access is difficult—therefore allowing for price discrimination.

So long as resale involves transaction costs, a limited degree of price discrimination can occur even in fairly competitive retail or industrial markets.⁷⁶ For example, in the hotel industry, some

customers pay higher prices than other customers for exactly the same hotel room. By charging a higher price to consumers with a higher willingness-to-pay for rooms (as determined by later date of purchase or the search method used by customers), the hotel can increase revenue. Suppliers also practice price discrimination when it costs more to supply one customer than another, yet the supplier charges both the same price. For example, all-you-can-eat buffets typically charge a single price to every consumer who purchases a meal. Some of those consumers eat larger portions and thus cost more to serve than customers who eat smaller meals. By charging each person the same price the restaurant hopes the small-portion customers will subsidize the large-portion customers and maintain the profits of this pricing scheme.

In the case of Internet policy, ISPs wish to begin price discriminating so they may charge content providers fees for a service they are currently providing but not charging for—access to their subscribers. Content providers would need to enter into contracts with each individual ISP to reach that ISP’s subscribers. Some websites (such as Yahoo, Facebook, and YouTube) derive a great deal of benefit from accessing Internet users. The price these websites pay ISPs to upload information onto the Internet is far below their willingness-to-pay for access to Internet users, meaning they extract a great deal of benefit (i.e., “surplus”) from the transaction, and they would pay more if needed. Other websites generate very little revenue (such as personal blogs, academic sites, and niche message boards), and those content providers may not be willing to pay much more than their current upload fees to access Internet users.

Without price discrimination, it is impossible to charge the large, revenue-generating websites differently than the personal blogs and academic sites that generate little revenue, and ISPs therefore do not charge for access to their subscribers. By not charging websites to access subscribers, ISPs are not tapping into content providers’ full willingness-to-pay. Some large websites benefit by paying the single market price when in fact they would be willing to pay much more.⁷⁷ This generates a benefit (i.e., a “consumer surplus”) for almost all content providers.⁷⁸


The availability of new technology called Deep Packet Inspection allows ISPs to determine more information about the content that flows over the Internet.

Previously impossible, ISPs could now charge different content providers different prices to reach broadband subscribers, which could have a tremendous impact on how the benefits created by the Internet are distributed.

If ISPs were allowed to charge different websites different prices to access their subscribers, content providers might have to pay a price closer to their full willingness-to-pay to access users. That is, by charging content providers the full amount they would be willing to pay to access subscribers, ISPs would be able to extract all the economic benefits from the transaction. As will be explained in Part Three, the process of charging each content provider exactly its willingness-to-pay is extremely delicate and difficult, and ISPs may not always be able to decipher a content provider's exact willingness-to-pay. The difference in willingness-to-pay across groups is a necessary condition for price discrimination, but alone it is insufficient to allow price discrimination.

DPI technology is the final component, making Internet market segmentation possible and therefore allowing price discrimination. DPI makes it possible for ISPs to identify the source and type of packets traveling over their network, and to treat them differently for pricing purposes. By using DPI to charge content providers their willingness-to-pay to access subscribers, and thereby extract surplus from content providers, ISPs would be able to transfer surplus from content providers to themselves. This would amount to a significant “wealth transfer” in the Internet market—a transfer of surplus from Internet content providers to Internet Service Providers. Ending the *de facto* net neutrality regime could force every content provider to pay each ISP to access the ISP's broadband subscribers. Content providers would then need to use part of their revenue to pay ISPs for access to subscribers.

Ultimately, there is a zero-sum game between surplus for ISPs and surplus for content providers. While surplus can grow with investment, at some point the ISPs and the content providers have to



split the surplus that exists. Because the surplus is finite, how it is distributed will determine winners and losers. Thus, while price discrimination can create more wealth for ISPs, it would decrease wealth for content providers. This would shift revenue from content providers to ISPs for the foreseeable future. On the other hand, the current *de facto* net neutrality system distributes wealth in favor of content providers, and the FCC rules would continue this distribution.

Figure 6 demonstrates a transfer of wealth through price discrimination. In the standard market graph (similar to Figure 2) the consumer and producer surplus are divided by the price line. That graph

represents the distribution of surplus for a single market price. If producers are able to charge each

consumer exactly their willingness-to-pay, they can extract the entire surplus from the market. This leaves the consumer with no surplus and the producer with the red triangle of surplus.

From a purely economic perspective, wealth transfers (*distributions* of surplus between market players) are considered neutral from an efficiency standpoint; market efficiency is only concerned with the *amount* of surplus generated by *all* market players.⁷⁹ In Figure 6, the quantity in the market equilibrium are unaffected by the wealth transfer, so that the market is still efficient and represents an optimal allocation of social resources. As long as the Internet information market produces the maximum possible level of economic surplus (defined as the sum of the producer and consumer surplus), then the market will be efficient. There is no reason to think *a priori* that the wealth should remain in the hands of the content providers as opposed to in the hands of the ISPs.

While they do not necessarily affect efficiency, wealth transfers are not irrelevant to policy debates—for example, there are important normative implications of wealth transfers from low-income people to higher-income people. In the case of the Internet, there is no obvious normative problem with a wealth transfer from content providers to ISPs. However, as will be discussed below, the wealth transfer enabled by price discrimination may exacerbate existing market failures and affect long-term incentives to invest in the market in a way that reduces surplus. Therefore, a comprehensive analysis of any policy should examine its effect on wealth distribution.

At its heart, net neutrality regulation is about who will get more surplus from the Internet market. Enforcing net neutrality keeps the surplus in the hands of the content providers. Eliminating it transfers the surplus to ISPs.

Changing wealth distribution will affect the abilities and incentives of the respective market players to invest in the portions of the Internet they own.

At its heart, net neutrality regulation is about who will get more surplus from the Internet market. Retaining net neutrality would keep more surplus in the hands of the content providers, and eliminating it would transfer some surplus into the hands of the ISPs. Changing wealth distribution would affect the ability and incentive of the respective market players to invest in the portions of the Internet they own. As will be explained in Part Three, allowing price discrimination would reduce the content providers' incentive to invest in the part of the Internet they own (content), while ISPs may have more incentive to invest in Internet infrastructure (the part of the Internet they own). This alignment of incentives would ultimately cause a problem because both content and infrastructure are subject to positive externalities. Thus, wealth transfers away from either are problematic because they would magnify already existing underinvestment. However, as will be discussed in Part Four, it is far easier for the government to make up the shortfall for infrastructure investment; protecting content providers' current surplus is the best policy option given the structure of the Internet market and the difficulty of directly subsidizing the creation of content.

Part Three

The Problem with Price Discrimination

Wealth transfers theoretically do not affect market efficiency. Nonetheless, there are several reasons the government should prefer to avoid a wealth transfer from content providers to ISPs in the Internet market. While the transfer of wealth may increase ISPs' ability and incentive to invest in creating new and stronger broadband lines (the Internet's infrastructure), it would decrease content providers' ability and incentive to invest in creating more information and websites on the Internet (the Internet's content).

Investments in both content and infrastructure are crucial if we wish to continue to grow the Internet, an important input into many industries vital to our economy—not to mention its importance to political participation, education, and the free exchange of information and ideas. Price discrimination would create a policy that incentivizes market investment in infrastructure at the expense of investment in content. To avoid underinvestment, the government would need to directly subsidize Internet content, but as will be explained in Part Four, content subsidization is an extremely difficult undertaking for a government, and the United States has had only limited success in this arena. Further, allowing ISPs to price discriminate does not ensure that ISPs will take the additional revenue and invest it back in the Internet infrastructure. In fact, much of the additional surplus would compensate for past decisions; only a small portion will incentivize future investments.

Further, perfect price discrimination is unlikely in the real world because it requires a company to know consumer preferences at a level of detail that is difficult to actually achieve. If ISPs imperfectly price discriminate, they will create inefficiencies in the Internet market, lowering the total surplus of the market and decreasing the value of the Internet.

A. Price Discrimination Will Alter Investment Incentives

Supporters of price discrimination argue that it will allow ISPs to collect more revenue, and therefore allow ISPs to invest in the Internet infrastructure, which they own. Although this is likely

true, supporters of price discrimination miss two key points. First, because the surplus generated by the Internet market is finite (as explained in Part Two), if more surplus were distributed to ISPs, ISPs would have more money to invest in infrastructure but content providers would have less money to invest in content. Second, price discrimination is an extremely inefficient tool to use to increase investment in Internet infrastructure. Price discrimination would mainly reward ISPs for past investment, as opposed to incentivizing future investment in infrastructure.

Increased Infrastructure Investment

As explained in Part One, there are two complementary goods that make up the Internet: infrastructure and content. For the most part, investments in these goods are made by several distinct groups of market players. ISPs, who own the broadband lines and most of Internet infrastructure, take part of their revenue (or raise capital) and invest in their infrastructure—for example, by expanding their broadband lines to new areas or increasing their speed. Content providers, who create websites and applications on the Internet, take part of their revenue (or raise capital) and invest back in their content—for example, by adding new information to their websites, creating new sites, creating new applications, or increasing the speed of current applications.

Supporters of price discrimination argue that if ISPs are able to charge content providers for access to the ISPs' subscribers, they would use that additional revenue to invest in their own networks. By directing revenue from content providers to ISPs, a price discrimination policy would increase investment in the hard-

There is a tradeoff between investment in content and investment in infrastructure. Price discrimination would increase incentives to invest in infrastructure, but reduce returns for investment in content.

It is vital to generate incentives to invest in both infrastructure and content to overcome the market failure of positive externalities.

wiring of the Internet and increase the adoption and quality of broadband in the United States. Increasing an ISP's rate of return on its investment in broadband lines by increasing the revenue earned on these existing lines would increase incentives to invest more in these lines.

If ISPs use their additional revenue to build out their broadband and wireless networks, the benefits to the U.S. economy could be substantial. The World Bank estimates that every 10% increase in broadband penetration in a developed country increases economic growth by 1.2%.⁸⁰ This is relative to a U.S. Gross Domestic Product growth rate of 2.8% over the last 20 years and a current broadband penetration rate of just under 50%. So, for example, if eliminating net neutrality were to increase broadband penetration even by 10% (from 50% to 60%), that penetration would increase the value of the U.S. GDP by \$289 billion per year. Assuming the 1.2% in economic growth is causally linked to increase in broadband penetration,⁸¹ if the additional revenue from abolishing net neutrality were in fact used by ISPs to improve access to broadband, the potential economic benefits could be substantial.⁸²

The quality of broadband Internet infrastructure in the United States is a hotly debated topic.⁸³ A recent report by the Berkman Center at Harvard University, *Next Generation Connectivity: A Review of Broadband Internet Transitions and Policy from Around the World*, finds that the United States is "a middle-of-the-pack performer" among wealthy countries in terms of broadband speed and uptake.⁸⁴ The report also finds that "[l]arge, long-term [capital] investments have played a role in some of the highest performing countries," and argues that other countries' open access policies appear to have increased those countries' level of competition and improved efficiency in the Internet service provision market.⁸⁵

However, the current quality of Internet access is direct evidence neither in favor of nor against net neutrality. The goal of any policy should be to maximize the value of the Internet, which means choosing a policy that addresses *both* the quality of broadband service and the quality of Internet content. Focusing exclusively on either of the two complementary goods may lead to overinvestment in one at the expense of underinvestment in the other, thereby reducing the total surplus in the market.

Decreased Investment in Content

While price discrimination would likely increase investment in infrastructure, it would also reduce revenue for content providers and thereby reduce their incentive to invest in Internet content.

As explained in Part Two, the surplus generated by the market is finite. If price discrimination distributes more surplus to ISPs, ISPs would have more money to invest in infrastructure but content providers would have less money to invest in content. There is a tradeoff involved in directing revenue from content to infrastructure. So, while price discrimination may improve the quality of Internet connections for new and existing subscribers, it may also decrease the quality of the Internet content that travels over those connections.

Internet content is extremely valuable and any alterations to incentives to invest in content may have serious economic consequences. For example, the net present value of surplus generated by the Internet content market is at least over \$300 billion dollars. This value represents the 2009 market capitalization for five of the largest Internet content providers⁸⁶—only a fraction of the total economic surplus generated by content providers. It does not include the market capitalizations of the thousands of other Internet firms, including some very large ones that are not publicly traded. It also does not include the fraction of market capitalization attributable to Internet activities at large conglomerate firms and the many other types of surplus that are generated by Internet content that is not represented in the value of Internet companies. While the actual surplus generated by Internet content is extremely hard to calculate, a floor of \$300 billion shows that the stakes are very high.

This impact on Internet investment is likely to be the largest consequence of the FCC's choice to develop net neutrality policy. Allowing price discrimination may increase investment in infrastructure, but at the expense of investment in content.

Even if price discrimination incentivizes investment in one part of the Internet, it will not necessarily increase investment in the Internet as a whole.

In order to maintain this large wealth benefit produced by Internet content, there must be a balance of incentives to invest in infrastructure and content. This balance is fragile and critical. Depending on how these incentives shift and to what degree they shift, they could affect the surplus generated by the Internet market. If investment in content is too high and investment in infrastructure is too low, then there would be plenty of information to access on the Internet but fewer users would be able to access it; content providers would then be unable to generate sufficient revenue, reducing investment and thereby decreasing the amount of information available. If investment in infrastructure were too high and investment in content is too low, then there would be many users on the Internet, but there would be less useful information for them to access.

The distribution of the surplus directly affects incentives and ability to invest in different parts of the Internet. This impact on Internet investment would likely be the largest consequence of the FCC's choice to enact net neutrality policy. It is insufficient for net neutrality opponents to simply argue that price discrimination incentivizes investment in one part of the Internet. The value of the

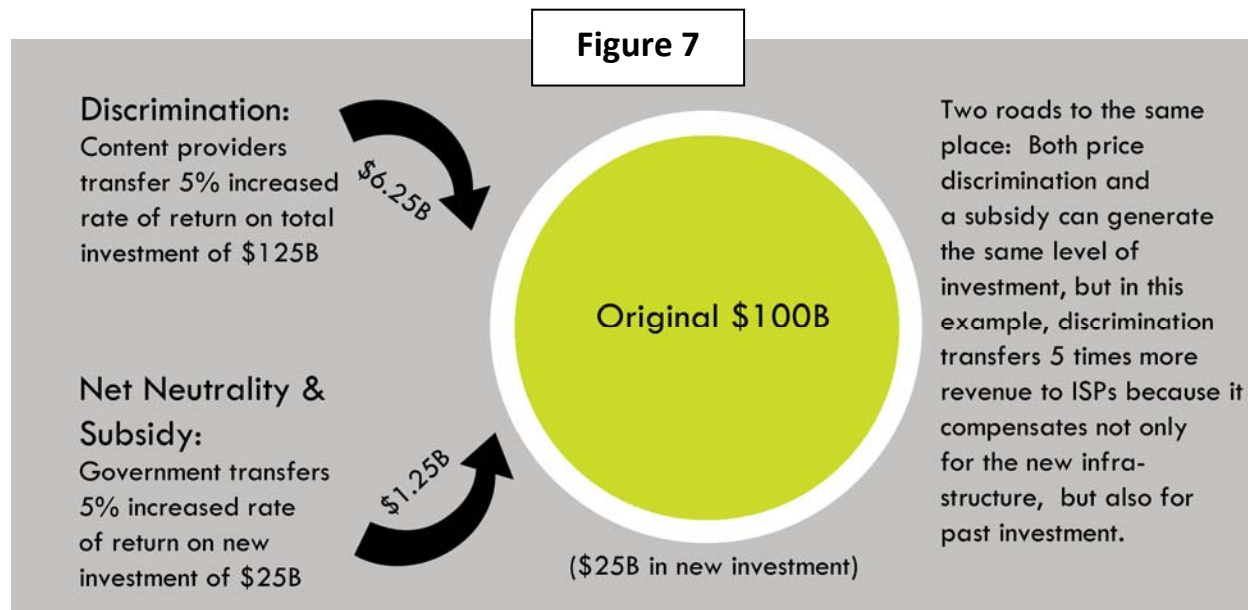
Internet is only partially related to the level of infrastructure provided by ISPs. Infinite amounts of infrastructure would not be valuable unless quality content were also available to access.

For this reason, regulators must consider the impact of net neutrality policy on the incentives to invest in infrastructure *and* content and how these two sectors together bring value to Internet subscribers.

An Extremely Blunt Tool

Price discrimination will not perfectly alter incentives to increase investment in new infrastructure. In fact, it is an extremely expensive tool to increase investment in infrastructure. Price discrimination in the Internet market would generate additional revenue for ISPs based on their past investment decisions, not based on consumer need for future broadband investment. Immediately after the elimination of net neutrality, all additional revenue generated from price discrimination would be based on existing Internet infrastructure. Only existing ISPs with completed infrastructure would be positioned to reach end-users and take advantage of the additional revenue that could be collected by charging content-developers for access to the end-users. Over time the additional revenue would incentivize ISPs to expand their network, but most of the benefits of eliminating net neutrality would go to reward existing investment. There are far cheaper methods to achieve and ensure investment in new infrastructure.

As Figure 7 below shows, only a small portion of additional revenue generated by ISPs from price discrimination would actually reward future decisions. In this example, price discrimination allows the ISP to earn a 5% greater rate of return on its existing investment of \$100 billion. That higher rate of return induces the ISP to invest in building out the network through a capital expenditure of \$25 billion, so that the ISP's total additional return from price discrimination is \$6.25 billion per year (5% of \$125 billion). Alternatively, a program of targeted government support could be used to incentivize the ISP to spend the \$25 billion by guaranteeing the additional 5% rate of return *on only that investment*, not the past decisions that are already locked in. The cost of this direct support would be \$1.25 billion per year. Price discrimination, then, transfers *five times* the amount of funds to ISPs while spurring the same level of investment. These funds are likely to be transferred to shareholders rather than being invested in expanding the broadband network.



In addition, because price discrimination is not targeted the way direct government support could be, the investment it generates would not necessarily reflect social priorities. For example, large established ISPs like AT&T and Comcast would be able to charge content providers substantial prices for access to their subscriber bases. Each of those two ISPs has over 14 million subscribers,⁸⁷ and content providers who could afford to pay the new ISP-specific access fees would likely make an effort to reach those millions of users. Those large ISPs would therefore earn this profit based solely on their existing broadband infrastructure. On the other hand, a company like Alaska Communications Systems serves approximately 55,000 customers,⁸⁸ and it would not be able to charge content providers a fee as high as those charged by AT&T and Comcast for access to its subscribers. That company would therefore not receive as much additional revenue and would not have as much revenue to invest in its lines—even though Alaska is particularly underserved for broadband Internet connection and is in need of new broadband infrastructure.⁸⁹ In this way, the wealth transfer effectuated by price discrimination will not match social needs for the creation of new broadband lines.

A large portion of the wealth transfer from content providers to ISPs would be essentially wasted because it would compensate for decisions that are already locked in, and most of the additional revenue would simply accrue on the basis of assets that the ISPs have already created.

Less Total Investment

Because competition is stronger in the Internet content market than in the market for broadband service, taking surplus from content providers and transferring it to ISPs may lead to a decrease in aggregate investment in the Internet. By most measures, market competition in broadband provision has declined over the last decade. The average consumer can now choose between two broadband providers, typically the local cable and phone companies.⁹⁰ And despite large profit margins, the industry barriers to entry are high, which reduces competition.⁹¹ Before entering the market, new firms must develop a full network of connections to local households at huge expense. These barriers to entry create significant natural monopolies in this market, making entry difficult. As an indicator of potential competition problems, the reported average monthly price of broadband access in the U.S. has increased from \$42.15 in 2003 to \$44.09 in 2007.⁹² Thus, ISPs do not need to be as vigilant about competition or use their investments to compete with potential new entrants into the market. Instead, they are free to use their additional revenue to generate proprietary content, invest in other parts of their business, pay dividends to shareholders, or reward managers with bonuses.⁹³

Price discrimination is an extremely inefficient tool to increase investment in Internet infrastructure because it will generate additional revenue for ISPs based on their past investment decisions, not based on consumer need for future broadband investment.

There are far better methods to achieve and ensure investment in new infrastructure.

Alternatively, there is strong reason to suspect that if wealth remains in the hands of Internet content providers, they will reinvest a greater share of their revenue into new and better forms of content. The market for Internet content is significantly more competitive than the market for Internet broadband access.⁹⁴ Internet content is constantly changing, with new product and players emerging at a furious pace—content providers must adapt (and invest in changing and adapting) to keep up with other content providers. The barriers to entry in the content market are low because it is relatively easy for a competitor to create a website over the network. If content providers do not improve and add websites and applications, they can lose users to new entrants into the market. If content providers attempt to pass revenue to their shareholders or employees instead of

investing in their content, other sites can quickly respond by improving their content and forcing out the content providers who do not reinvest their profits. The value created for content providers' by past investment in their websites and applications will evaporate quickly if they do not continuously reinvest their revenues.

The effects of price discrimination on investment are complex, but there are strong reasons to think that, while it will lead to increased investment from ISPs, it will come at the expense of investment in content that would have generated significant positive externalities. Further, the same level of investment from ISPs can be achieved through direct government support, which can be better targeted to ensure that broadband investment tracks social needs. Finally, given the relative difference in competition between the two markets, transferring surplus from content providers to ISPs may result in decreased aggregate investment in the Internet—a significant problem given the existence of positive externalities. On balance, then, the overall investment effects of price discrimination are likely to reduce the total surplus generated by the Internet over the long run, compared to a policy of net neutrality.

B. Imperfect Price Discrimination Could Reduce Internet Surplus

While perfect price discrimination in the Internet market raises several important concerns, there is another set of issues created by the practical difficulties of actually implementing price discrimination in the real world. If price discrimination is not carried out perfectly, it could cause inefficiencies if ISPs fail to charge the “right” prices and drive some valuable content providers out of the market. This loss of valuable content would then lead to a decrease in surplus in the Internet market—a loss for society as a whole. In addition to these informational problems, ISPs also face perverse incentives that may lead to a decrease in surplus. Finally, price discrimination will also reduce surplus by extracting revenue from the existing Internet market and using it to set up a price discrimination regime.

Difficulty in Estimating Willingness-to-Pay

Execution of perfect price discrimination requires a great deal of information, as ISPs must be able to accurately ascertain each content provider's willingness-to-pay⁹⁵ for access to the ISP's subscribers as potential consumers, and charge each content provider exactly that price. Only then can an ISP gain the entire surplus from the transaction. If price discrimination is imperfect, as it would likely be, then the wealth transfer from content providers to ISPs could alter the quantity of goods or services that change hands in the market and create market inefficiencies.

The information ISPs need to properly price access to their network is difficult to collect. Ideally, from the ISP's perspective content providers would indicate to the ISPs what they would be willing to pay to access end-users, and ISPs could charge content developers exactly that price. Of course, content providers have an incentive to understate their true willingness-to-pay in order to minimize the price. It would therefore be difficult for ISPs to determine each content provider's exact willingness-to-pay.

Unless ISPs can develop a creative pricing scheme that encourages content providers to voluntarily reveal their true willingness-to-pay (an unlikely scenario), ISP price discrimination would likely be far from perfect. If ISPs charge prices that do not exactly match each content provider's willingness-to-pay, either ISPs would lose part of the surplus they could have gained, or content providers will exit the market. If an ISP were to charge content providers too little, the ISP would not maximize its surplus from the transaction, leaving the rest of the surplus with the content provider. For example, if a content provider like YouTube were willing to pay \$10,000 per month to access an ISP like Time Warner's subscribers, but Time Warner only charged YouTube \$8,000, Time Warner loses out on

\$2,000 of potential revenue. This result would not be inefficient from the market's perspective, but would represent a loss from the ISP's point of view.

If an ISP charged a content provider too much, there would be negative consequences because it might force those content providers to exit the market, which would reduce the economic surplus generated by the Internet. For example, if an ISP like AT&T charged a content provider like Orbitz a price higher than Orbitz's willingness-to-pay, Orbitz may choose not to access AT&T's subscribers, and opt to contract only with other ISPs that charge a lower price. AT&T should, then, simply lower its price in order to capture surplus from Orbitz's interactions with AT&T subscribers. There are several reasons why this may not happen however. AT&T cannot know if Orbitz has revealed its true willingness-to-pay, or has simply engaged in a negotiating strategy of signaling a lower reservation price that it actually has—the two parties may end up in a “game of chicken” that would result in no bargain. Alternatively, AT&T may have misclassified Orbitz into a higher value category that it should be, and it may be difficult for AT&T to undo the categorization without opening up a large number of other contracts. If any type of problem like this were to persist, and Orbitz and AT&T could not arrive at a deal, the failure to engage in the transaction would result in a loss of revenue for AT&T, Orbitz, and AT&T subscribers (who would be given fewer choices of content). From an investment perspective, the potential for these types of breakdowns increases the uncertainty of returns on capital invested in the Internet, and therefore would reduce the rate of investment

Failure to engage in efficient contracts could also impact the ability of Internet users to share content freely, destroying a significant portion of the Internet's special value. For example, if a company like Verizon were to charge a price above a content provider like the *Chicago Tribune's* willingness-to-pay, the newspaper would exit from that ISP's market. When subscribers at other ISPs saw content at the *Tribune* website they wish to share, they would be unable to provide links to Verizon's customers. The result is a drag on the free exchange of information between customers at different ISPs—a situation that is currently unimaginable for Internet users.

Scenarios like these would create Internet fiefdoms and drastically reduce the value of the Internet by destroying its unrestricted end-to-end dynamic architecture. With less content, less information, and fewer websites, the network could become less valuable for users; and some users may leave the network, making it less valuable for other content providers and other users.

In an environment with network effects, this loss of surplus is potentially magnified. Reducing the number of users on the Internet would lower the value of the Internet for those who remained; and those other users may then also choose to leave the Internet. When users leave the Internet, it creates losses of surplus as those users are no longer engaging in surplus-creating transactions. In addition, loss of users would also reduce the amount of information posted to the Internet: this would lower the value of the Internet for users who remain. Because of the reduction in value of the Internet, more users might leave; the more users who unplugged from the network, the larger the loss of surplus. This negative feedback would reduce the network's value to both content providers and users.

Incorrect Incentives

Because ISPs see all of the revenue they capture from price discrimination, but face only a portion of losses if they fail to price discriminate perfectly, they would not have the correct incentives to reduce price discrimination errors to the social optimum. ISPs would prefer to have a much larger piece of a smaller Internet surplus pie than a much smaller piece of a larger pie. As long as ISPs' revenue is maximized, ISPs would not care about losses in other parts of the Internet market.

ISPs are not concerned with the well-being of content providers, subscribers of other ISPs, or the Internet as a whole. If ISPs could maximize their own surplus, they would not care if the surplus from the Internet as a whole was reduced. If the loss of content and subscribers in the market were offset by the increase in revenue to the ISP, ISPs would choose to charge content providers despite the broader impact on the value of the network as a whole.

For example, an ISP like Verizon could choose to offer take-it-or-leave-it contracts to large content providers of \$10,000 per month for access its subscribers. If a company like Wikipedia had a willingness-to-pay lower than \$10,000, but companies like Yahoo and eBay had a willingness-to-pay that is higher, Wikipedia will not transact with Verizon but the other sites will. So long as the value of any subscribers lost due to the loss of Wikipedia access does not add up to \$20,000, Verizon will come out ahead. In this case, Verizon is better off, even if Wikipedia and all users were worse off, so Verizon would continue to charge this price that would be too high for Wikipedia to pay.⁹⁶

Because of these imperfect incentives, ISPs are unlikely to consider the impact of their price discrimination policies on a sizeable fraction of the Internet market. ISPs would therefore err on side of extracting more money from content providers by setting higher prices even if it were to result in some content providers leaving the network and a lower value for the Internet as a whole. This makes it possible that their pricing schemes would lower the value of the Internet while increasing revenues for ISPs.

Costs of Switching Regimes

Price discrimination may reduce the surplus in yet another way—by taking surplus out of the Internet market and using it for transaction costs. Switching from a system of net neutrality to price discrimination will involve significant transaction costs. ISPs will need to set up a system to employ price discrimination, implementing DPI and setting up systems to charge different content providers fees to access their subscribers. Content providers will need to spend time figuring out which ISPs they need to contract with, as well as time making decisions about whether to enter into contracts with them and at what price. Content providers will need to evaluate the conditions of the ISPs' contracts and spend time entering into and executing these contracts. Transaction costs are essentially wasted money; they generate no benefits for ISPs, subscribers, or content developers. The time spent on evaluating and entering into agreements would move money out of the Internet market—revenue that could have been used for investment in Internet content and infrastructure—into the deadweight of transaction costs. In this way, these transaction costs will deplete some of the surpluses of the Internet market. However, as mentioned above, depending on how much revenue ISPs think they could extract from content providers, they may still prefer paying the transaction costs of switching to price discrimination.

Uncertainty

Additionally, the impact of changing over to a price discrimination system involve uncertainty, another cost. Some argue that policymakers should not enact net neutrality regulations because it will have an uncertain impact on consumers.⁹⁷ However, the opposite argument applies more forcefully—because net neutrality represents the *status quo*, there are greater uncertainties involved in switching to an untested price discrimination regime. The Internet market in the United states has never functioned under price discrimination, and switching to a new system involves a host of unknowns, including whether ISPs can perfectly price discriminate, whether ISPs would invest their additional revenue in infrastructure, and whether eliminating net neutrality would increase competition in broadband. The Internet has always functioned under a *de facto* net neutrality rule. Any estimates of the benefits or costs to consumers of abolishing net neutrality are

therefore somewhat speculative.⁹⁸ Given current understanding, changing the *status quo* involves more uncertainty than maintaining it.⁹⁹

For these reasons, price discrimination is a risky policy option. There is significant risk that it would decrease the total surplus in the Internet market, especially if ISPs did not have enough information or enough incentive to perfectly price discriminate. Even if ISPs were able to perfectly price discriminate, such a policy would at best incentivize investment in infrastructure at the expense of investment in content. As Part Four will show, there are far better and more efficient ways to increase investment in Internet infrastructure.

Zero Price

Given the problems associated with price discrimination, it might be argued that the alternative of a standard single price for all content providers to access an ISP's users (set according to the number of users on an ISP's network) would be the best pricing mechanism. However, such a pricing model would have many of the problems associated with price discrimination and would be *guaranteed* to create some of the worse consequences of imperfect price discrimination.

Any positive price for content providers to access an ISP's subscribers will lead to a loss of surplus for content providers. While ISPs would see an increase in their rate of return, it would be at the expense of content providers, leading to the investment problems discussed above: any increase in broadband investment would be less than, or at most equal to, the reduction in investment in Internet content. At best this would result in no social loss, but would be more likely to cause reductions in the aggregate surplus created by the Internet.

Further, just as imperfectly executed price discrimination could cause content providers to exit from some markets and decrease market surplus, a single price for access to subscribers would also cause the exit of those content provider that have a willingness-to-pay below that price. This exit would destroy some of the surplus enjoyed by both content providers and the end-users, as well as the positive externalities associated with that content.¹⁰⁰ In this way, a zero price may lead to the most beneficial results for the Internet market.

While in a perfect market, price discrimination represents a neutral wealth transfer, and does not affect the total surplus created by the market, there are a variety of reasons to be concerned about the consequences of price discrimination in the Internet market. Even if done perfectly, price discrimination would at best be a zero-sum game, and increased revenues for ISPs would come out of the pockets of content providers. Because of the relative competitiveness of the markets for broadband and content, there is reason to believe that this transfer would result in less investment in the Internet, because broadband owners would be more likely to take increased revenue and distribute it to shareholders rather than reinvest it. Further, much of the increased rate of return generated by price discrimination would be for past decisions; while increased rates of return would create incentives for ISPs to expand their networks, the same results could be achieved through a system of targeted government support. All of these problems exist even if ISPs price discriminate perfectly—errors in price discrimination will lead to even greater losses in surplus as content providers exit the market. Given all the downsides associated with price discrimination, alternative policies to incentivize infrastructure investment are likely to have more efficient results.

Part Four

Improving Infrastructure Through Direct Government Support

As explained in Part One, given the market failure associated with information in the Internet access market and the resulting under-provision of information, government may be justified in attempting to correct this failure. There are a number of ways governments can attempt to correct for an externality in any market, from the creation of new market mechanisms or forms of property (such as patent rights or emissions allowances), to direct regulation, to investment subsidies.

One example of government attempts to correct for positive externalities involves the research and development of new pharmaceuticals. A company can invest in research and generate a new drug that provides a great deal of benefits to society. However, other companies, if allowed, will copy that drug and sell it at a lower price (because their cost of producing the drug does not include the cost of research and development to invent it). This leaves the original company unable to capture all the benefits from its investment in the invention and disincentivizes drug companies from investing in research and development. The U.S. government corrects this market failure by granting patent protection for a limited time to new drugs. The opportunity to have a short-term monopoly and collect monopoly profits grants pharmaceutical firms an incentive to invest in research and development. In this way, federal patent laws correct the market failure by *internalizing* some portion of the positive externality in the pharmaceutical market.

A government could also correct a positive externality by subsidizing producers to incentivize them to increase production to a socially efficient level. In the context of pharmaceuticals, the government funds a great deal of primary research in the area of biology and genetics that continues to inform the development of new drugs. Operating alone, the market would not support this type of research, and so the government steps in to ensure that the market produces this good (primary research) with its large potential positive externalities.

In the Internet market, the federal government should establish a policy that produces results as similar as possible to an efficient market in equilibrium, by correcting for the effects of positive externalities. Because of the complementary nature of Internet content and access, simply having an FCC policy that allows price discrimination and increases incentives for ISPs to invest in

network infrastructure without somehow increasing investment in content would be insufficient to correct the market failure associated with information provision on the Internet. Content and infrastructure must be broadened in concert to fully address this issue.

If a large quantity of resources were devoted to Internet infrastructure and no effort were made to expand the content, we might eventually end up with fast, ubiquitous access to low-quality content. On the other hand, simply enshrining net neutrality and incentivizing content providers to invest in content without increasing investment in infrastructure would create the opposite problem: a wide range of content that would be suboptimally available.¹⁰¹ For this reason, both infrastructure and content must be addressed by network policy.

As explained in Part Three, when evaluating its policy options in the Internet market, government should attempt to correct the market failures to maintain efficient levels of investment in both parts of the Internet, infrastructure and content. Thus, the government is faced with two

choices: establish a pricing policy that incentivizes investment in infrastructure, and try to directly subsidize the content side of the market; or establish a pricing policy that increases incentives for investment in content, and create a separate policy directly subsidizing Internet infrastructure. Given the relative ease of subsidizing infrastructure, and the difficulty of subsidizing content, the first policy—of establishing pricing policy that favors content and then subsidizing infrastructure directly—would likely lead to the best results.

Because of the complementary nature of Internet content and access, simply having an FCC policy that allows price discrimination and incentivizes ISPs to invest in network infrastructure without somehow increasing investment in content is insufficient.

Policy must broaden content and infrastructure in concert to fully address systematic underinvestment.

A. Relative Ease in Subsidizing Infrastructure

Internet infrastructure is not a typical market good. As a “public good,” infrastructure construction in general is often seen as the province of the government.¹⁰² Public goods are nonrival (meaning one person’s consumption does not limit another’s) and non-excludable (no one can be excluded from consuming the good). Private goods, on the other hand, are both excludable and rival, meaning only one person can consume them and individuals who do not pay for the good cannot consume it. Public goods are subject to a free-rider problem, in which no one individual wishes to pay for the production of the good because she can enjoy the same benefits if someone else pays for it. National defense and broadcast television are classic examples of public goods.

Goods such as Internet access and satellite television are most correctly classified as “club goods” because they are nonrival but excludable.¹⁰³ These goods are often provided by natural monopolies because the huge costs of creating an excludable access network serve as a barrier to entry for potential competitors. The costs of setting up a network have been estimated to be between 80%-90% of the total costs of providing broadband access.¹⁰⁴ These huge upfront (or fixed¹⁰⁵) costs prevent new ISPs from entering the market and competing against the incumbents. As is the case with public goods, the market does not always correctly provide club goods. Although not subject to the free-rider problem, the providers of club goods can often take advantage of the structure of the market to exercise market power and charge monopoly rents.

Governments often respond to these natural monopolies by setting up systems of price regulation in those markets to ensure monopolists do not charge users exorbitant prices. For example, many

electric utilities have prices set by utility boards. These boards typically set prices so that the electric utilities earn enough to cover their costs and provide reasonable returns on capital to their shareholders.

In addition to directly regulating natural monopolies, the United States has a long history of subsidizing infrastructure construction. These projects build and maintain public goods and extend access to club goods with the goals of increasing national defense,¹⁰⁶ creating jobs,¹⁰⁷ and improving long-run economic growth.¹⁰⁸ There is a considerable cottage industry evaluating the efficacy of these projects in accomplishing their stated goals. From this literature, it is clear that the federal government has had considerable successful experience participating in infrastructure projects.

The Interstate Highway System (formally the Dwight D. Eisenhower National System of Interstate and Defense Highways) is perhaps the most famous example of a public infrastructure project carried out at the federal level. In 1966, the United States government rolled out this program to connect major cities and border crossings with freeways in order to improve safety, reduce congestion, create jobs, and improve national defense in case of atomic war.¹⁰⁹ The economic impact of the highway system is widely debated. Some studies have shown that every \$1 spent on highway construction has created \$6 in returns;¹¹⁰ other observers argue that some of the economic growth created by the investment was simply redistributed from other areas.¹¹¹ Regardless of the exact economic benefit, the highway system is undoubtedly been a success. These highways serve nearly all major U.S. cities, and virtually all goods and services distributed in the nation travel over them at some point. Residents commonly use urban Interstates to travel to their places of work, and for the vast majority of their long-distance travel, whether for vacation or business. The entire system has over 46,000 miles of expressway, making it both the largest highway system in the world and the largest public works project in history.¹¹² It has had an enormous impact on the nation, “contribut[ing] mightily to the economic growth and quality of life” and allowing the United States to become a “nation on wheels” with “a freedom of mobility that is unrivaled anywhere in the world.”¹¹³

Because private markets tend to underprovide public goods, like Internet infrastructure, government can play a role in correcting these failures through direct investment.

The United States government has historically been successful in subsidizing public infrastructure, yet has had greater difficulty subsidizing content.

The United States government has also successfully subsidized club good projects. The Rural Electrification Administration (REA) is considered one of the New Deal’s most successful and cost effective programs. In 1935, electricity was not widely available in rural areas; only 10% of rural households had access while 90% of urban households were electrified. The REA was charged with addressing this imbalance. Widespread electricity would allow rural families to store meat without spoilage, complete indoor tasks during the evening, and increase the market for electrical appliances such as stoves and dishwashers.¹¹⁴ The agency offered financing on generous terms to rural co-operatives that sought to bring electricity to their communities. Incumbent electricity providers opposed the REA on both political and economic grounds, arguing that the federal government should not enter their market and worrying that REA-funded co-ops might compete with them on the edges of their network.¹¹⁵ By 1939, the REA managed to connect 288,000 households through 417 co-operatives, and the fraction of rural households that had access to electricity increased to 25%.¹¹⁶ The REA’s subsidization also helped incumbent electric companies recognize the profitability of serving rural customers despite the upfront costs of extending the

electric grid. The REA's program sped the process of widening access to electricity by not only funding formation of co-ops, but also by encouraging electric companies to act before co-ops formed in their region.

Moreover, for the better part of a century the federal government has been subsidizing communications infrastructure, and, more recently, is subsidizing infrastructure for broadband connection to the Internet. The Communications Act of 1934 implemented federal policy goals commonly called "universal service" in order "to make available . . . to all the people of the United States a rapid, efficient, Nation-wide, and world-wide wire and radio communication service with adequate facilities at reasonable charges."¹¹⁷ The federal government successfully achieved universal telephone service through vigorous FCC regulation: in 1984, the FCC and state governments began offering direct subsidies to low-income users, with the Lifeline program covering telephone service and the Link-up program covering installation.¹¹⁸

The Telecommunications Act of 1996 increased federal aid to extend to universal Internet access. The Act directed the FCC to expand coverage according to seven "universal service principles," including affordability, universal access to "advanced" telecommunications, and reasonable equality in cost between urban and rural areas.¹¹⁹ The FCC does this through the Universal Service Fund, which administers four separate subsidy programs: a high-cost program, which provides subsidies to carriers in order to keep rates low in rural or other high-cost areas; the low-income program, which includes Lifeline and Link-up; E-rate, which provides low-cost telecommunications for schools and libraries; and the Rural Health Care program, which does the same for rural hospitals and health providers.¹²⁰ As of 2004, the universal service programs have helped telephone service penetrate 95% of the market.¹²¹

More recently, the American Recovery and Reinvestment Act of 2009 allocated \$7.2 billion for expanding broadband access.¹²² The first \$4.7 billion was directed to the Broadband Technology Opportunities Program (BTOP), administered by the Department of Commerce, in order to provide broadband access to underserved communities and public institutions, as well as to stimulate broadband demand in general.¹²³ The BTOP allows local groups to apply for funding for "last mile projects" to bring broadband Internet service to census blocks where a majority is not served by broadband-speed Internet or have Internet speeds too slow to be considered broadband (less than 3 megabytes per second).¹²⁴

The remainder of the funds established a \$2.5 billion Broadband Initiatives Program, administered by the Department of Agriculture, which specifically aims to expand broadband access in rural areas.¹²⁵ The stimulus also authorizes the FCC to create a "National Broadband Plan" which ensures universal broadband access.¹²⁶ Several draft FCC proposals include regulatory reforms, overhauls to existing subsidy programs, and redirecting money from the Universal Service Fund.¹²⁷ Much like the REA, the aim of the government's communications programs is to subsidize communications infrastructure in underserved and rural regions and create jobs during an economic downturn.

These programs have already distributed approximately \$200 million in grants and loan to build Internet infrastructure including "\$33.5 million grant with an additional \$8.8 million in matching funds to deploy a 260-mile regional fiber-optic ring to deliver gigabit broadband speeds, reliability, affordability, and abundant interconnection points for last mile service in the North Georgia foothills," and "\$20.6 million grant with an additional \$5.1 million in matching funds to add 140 miles of backbone network and 219 miles of middle mile spurs to existing network, enabling the delivery of at least 10 Mbps service to more than 220 existing anchor institution customers in rural and underserved areas of the state."¹²⁸ The National Economic Council found that these programs would create jobs and "could create hundreds of thousands of jobs over a four-year period by

stimulating new businesses, market transactions, and innovative industries in previously underserved areas.”¹²⁹

These are just a handful of examples of federal government subsidization of infrastructure. There is nothing new about the government directly subsidizing infrastructure. There are strong economic rationales for using direct subsidies to ensure that infrastructure or other public goods projects are undertaken: because private markets tend to underprovide club goods, like Internet infrastructure, there is a role for government to take appropriate steps to correct these failures through direct investment. And, the United States government has historically been successful in subsidizing public infrastructure.

B. Difficulty in Subsidizing Content

The United States government has sought to subsidize content in several different ways with mixed results. The quality of content is very subjective, meaning that broad-based efforts to support content in a systematic way are likely to come up against opposition from people who do not like certain types of content.¹³⁰ The government is not well positioned to make these allocation decisions and is unlikely to be successful at selecting types of content on the scale necessary to make up for lost revenue due to price discrimination.

The Corporation for Public Broadcasting is an example of an effort by the United States government to support specific types of high-value content. For 70 years the federal government has funded radio and television stations to support educational programming.¹³¹ Approximately 20% of the Public Broadcasting System’s funding comes from the federal government, primarily through the Corporation for Public Broadcasting. The balance is provided by state and local governments and private donations.¹³² Through this effort, the U.S. seeks to provide types of programming underprovided by market stations to small viewer groups. These groups are usually too small or insufficiently wealthy to appeal to advertising that supports market stations. These are typically programs that are believed to provide a high social value, but are difficult for advertisers to monetize, such as programming for young children or specific ethnic groups, and cultural and historic programming. The federal government’s involvement in funding content for underserved markets is not without controversy. Some groups complain that programming does not represent their own viewpoint or represents views with which they strongly disagree.¹³³ Other opponents argue that the government should not be in the business of supporting content at all when a market exists that could serve those niche clientele.¹³⁴ They argue that the market remains un-served precisely because the government became involved in providing that for free. This debate reignites each time Congress re-approves funding for public broadcasting.

The United States has also sought to support content through the National Endowment for the Arts (NEA). Congress established the NEA in 1965 as an independent agency of the federal government

“dedicated to supporting excellence in the arts, both new and established; bringing the arts to all Americans; and providing leadership in arts education.” The NEA’s role was to increase the availability of art to citizens of the United States; it is now the nation’s largest funder of the arts. The NEA provides funding to artists chosen through a peer review process, in which fellow artists select

Content quality is subjective—efforts to support content in a systematic way are likely to meet opposition from individuals who do not agree with certain types of content.

The government is unlikely to be successful in selecting content on the scale necessary to compensate for lost revenue due to price discrimination.

recipients based on the promise of their work. The funding of art generally, and the funding of controversial artists in particular, has been extremely contentious. President Reagan attempted to eliminate funding for the NEA shortly after taking office.¹³⁵ In 1990, four performance artists whose proposals were approved by the peer review system at the NEA were stripped of funding by the Chairman of the NEA due to their controversial nature. The artists sued and were eventually awarded compensation equal to the amount of their lost grants.¹³⁶

The controversy surrounding the NEA is a window into the difficulty of subsidizing content and other subjective endeavors. Any time the government involves itself in a content market, its actions are open to criticism on the basis of the type of content it chooses to subsidize, or the type of content it chooses not to subsidize. When government provides content, it can never please all citizens—as some will almost always disagree with the views funded by the government. The funding problems experienced by the NEA would also plague any widespread effort to support content on the Internet.

History has shown that the federal government has had difficulties in subsidizing content, but that it can successfully subsidize infrastructure at a high benefit to society. Moreover, the United States is already actively subsidizing broadband Internet infrastructure. Abandoning the net neutrality framework represents a significant transfer of resources from the content side to the infrastructure side of the market, which will only serve to increase the information externality. By continuing the government's subsidization of broadband access and protecting the revenues earned by content providers from price discrimination by ISPs, the government can address the information externality without being drawn into the subsidization of content, which has proven so difficult in the past.

Thus, when faced with the choice of how to correct for the externalities in the Internet market, government must take into consideration its whole range of policy options. Given the government's historic success in subsidizing infrastructure and difficult in subsidizing content, it makes the most sense for government to correct the externalities by instituting net neutrality—a pricing policy that incentivizes market players to invest in content—and then directly subsidizing investments in infrastructure.

Part Five

Fast Lanes to Nowhere: The Risks of Prioritization

“Prioritization” involves ISPs placing content into “tiers,” giving priority access of their lines to content providers in higher tiers in exchange for a fee, and leaving other content providers’ websites less easily accessible. Whether ISPs should be allowed to prioritize is one of the most controversial issues in the net neutrality debate.

Under the most commonly proposed prioritization scheme, an ISP would choose to prioritize based on the origin of packets. For example, an ISP like Verizon could charge a content provider like *The New York Times* a fee to use a “fast lane” broadband access to send its packets to Verizon’s subscribers; other smaller sites like a local newspaper, that would not or could not pay for this faster access, would be left on the standard access (the “slow lane”). Under most conceptions of prioritization, ISPs would charge a fee for fast lane access but continue to provide free access to the standard (i.e. slow) lane. Generally, when the network is not congested, users would likely experience no discernible difference when accessing sites on different “lanes.” However, if the network were congested, Verizon could delay or drop packets from low priority websites like local newspapers so that it can allocate more bandwidth to websites paying for high priority access, like *The New York Times*. The high priority websites would load for Verizon’s subscribers at the same speed as during times of no congestion. Because Verizon’s subscribers could find it more difficult to access low priority websites, those sites would be less valuable to those subscribers.

ISPs could also choose to prioritize based on type of packet—for example, by charging content providers to place live streaming video on the fast lane, and allowing standard Internet web pages to remain on the slow lane. If the network were to become congested, then the ISP would deliver each packet of streaming video as it arrives, but slot data packets from standard websites between video packets so that the video application would run unimpeded. Depending on the amount of congestion, end-users may experience degraded performance when accessing lower priority applications, thereby reducing the value of those applications for those users.

Prioritization is a form of an economic concept termed “product differentiation.” A seller practices product differentiation when it provides two similar goods with differences in key features or

minor details at different prices. Economists typically endorse product differentiation because it allows buyers to choose among different options and it increases the chance of consumers finding a good that more accurately meets their needs. ISPs employing prioritization schemes and developing Internet fast lanes and then selling that access at a different price than standard access would be an example of sellers practicing product differentiation.

Some observers argue that prioritization's costs likely outweigh its benefits.¹³⁷ Others suggest that prioritization is crucial for network management and should be allowed under the FCC regulation.¹³⁸ The FCC's proposed

net neutrality rules would clearly prohibit pure price discrimination, preventing an ISP from charging different content providers different prices for using the exact same service—in this case the ISP's last mile broadband lines to its subscribers. As proposed, the FCC rules also appear to ban the practice of "prioritization" by ISPs.

The nondiscrimination rule would prohibit an ISP from treating any content, application, or service in a "discriminatory" manner, subject to reasonable network management. The proposal notes that this rule prevents an ISP from "charg[ing] a content, application, or service provider for *enhanced or prioritized* access to the subscribers of the broadband ISP."¹³⁹ The next sentence of the proposal states that rule "would not prevent a broadband Internet access service provider from charging *subscribers* different prices for different services." This discussion within the regulation implies that fast lane access for content providers would be prevented. The regulation appears to allow ISPs to charge their subscribers higher prices for better services, but prevents ISPs from charging content providers higher prices for better access to the ISPs' subscribers.

Prioritization can be a method of implementing price discrimination. The different lanes of traffic can be used to sort content providers by willingness-to-pay, and ISPs can use this information to price discriminate.

If that happens, all of the problems of price discrimination will also be created by prioritization.

Prioritization, although not necessarily harmful to the Internet market, can have large negative impacts on the value of the Internet and should therefore be viewed as a risk-laden policy option. First, ISPs could use prioritization as an opportunity to introduce price discrimination. In theory, this opportunity may not be a problem, but for the reasons described in Part Two and Part Three, price discrimination can be harmful in this particular market.

Second, when executing prioritization, ISPs face perverse incentives that can potentially lead to a loss of surplus in the market, which could lower the value of the Internet to society. ISPs consider only the value of their own network and not the value of the Internet as a whole when developing their network management and prioritization schemes. These incentives could lead to pricing that could increase the value of the ISPs' private networks, while reducing the value of the Internet for society as a whole.

A. Price Discrimination in Disguise

Most net neutrality supporters argue that prioritization will create the same hazards as price discrimination. Under such a scheme, only bigger and richer companies would be able to afford fast lane prices, which is "anathema to a culture of innovation,"¹⁴⁰ even if the only barrier to fast lane access were price.¹⁴¹ Smaller and less profitable companies, like startups, would be unable to pay for access to the fast lane; users would have a harder time accessing those sites as ISPs would be more likely to delay or drop packets from those sites during times of network congestion. Users

would then turn to competitors' sites that were on the fast lane, and thus faster to access, thereby decreasing the number of users accessing startups' sites on the slow lane. Depending on how many users drop off those startups sites, the smaller companies could be forced out of business, or perhaps never even launch their businesses, thereby impeding innovation on the Internet.

Most net neutrality supporters also argue that fast lanes allow ISPs to treat packets from less profitable sites in a less favorable and arguably "discriminatory" manner, taking the power to choose content out of the hands of subscribers and placing it into the hands of ISPs.¹⁴² Such a system may make the Internet look more like cable television—where network owners make content decisions on the basis of perceived consumer demand—and eliminate the current unrestricted end-to-end architecture, in which end-users make content choices directly.

Supporters of price discrimination and prioritization argue that allowing fast lanes will increase investment in Internet infrastructure.¹⁴³ Similar to the pro-price discrimination argument, they argue that ISPs will use revenue from fast lanes to invest in their broadband lines, creating more and faster Internet access across the country. They also argue that without the existence of fast lanes, new businesses like telemedicine will be unable to flourish.

The debate, however, must look beyond whether textbook product differentiation in a theoretical market may be beneficial or harmful. Prioritization *as executed* by ISPs within the Internet market could have detrimental consequences. One of the most important arguments against prioritization is its ability to quickly turn into price discrimination. Economists have recognized this phenomenon, which occurs when sellers price the higher-quality product at a price that is higher than justified.¹⁴⁴

Under standard product differentiation, the seller would offer a slightly higher-quality product at a slightly higher price, and offer a standard product for a slightly lower price. This difference between the high price and low price should be the seller's marginal cost to produce the higher quality product. Under such a pricing scheme, the surplus generated by the Internet could increase: content providers would be able to choose from among different speed options and increase their chances of finding an access tier that would more clearly meet their needs.

But sellers could choose to price the higher-quality good at a price *higher* than the marginal cost to make the good of higher quality. In the Internet market, ISPs could price access to the fast lane at a price higher than the ISPs' cost difference to provide that higher level of service. Under proper product differentiation, the two products must be different enough to justify the higher price. Any difference in price between the fast lane and slow lane that is not proportional to the difference in cost between providing fast lane and slow lane access would have an element of price discrimination: instead of offering two different products at two different prices, the seller would in essence be offering almost the same product at two different prices, a form of price discrimination. In such a case, ISPs would basically be sorting content providers by willingness-to-pay and then extracting as much revenue as possible. The price difference would not be justified by the cost of providing different products, but instead would result in a wealth transfer from content providers to ISPs.

This transfer would allow for all the potential negative effects of price discrimination discussed in Part Three. While the transfer of wealth may increase ISPs' ability and incentive to invest in the Internet's infrastructure, it would decrease content providers' ability and incentive to invest in the Internet's content. To avoid underinvestment in content, the government would need to directly subsidize Internet content; and as explained in Part Four, content subsidization is a relatively difficult undertaking for the government. Further, allowing ISPs to extract this additional revenue would not guarantee that ISPs will invest that revenue back into Internet infrastructure.

Finally, as we will see in the following section, the availability of prioritization will create some perverse incentives that may motivate ISPs to engage in these types of “price discrimination in disguise” scenarios. These perverse incentives will eventually create the conditions for market inefficiencies that may decrease the entire value of the Internet market, leading to a suboptimal economic outcome.

B. Prioritization May Reduce Surplus

There is no guarantee that ISPs will use prioritization to maximize the surplus of the Internet market. There is, however, strong reason to think that ISPs may engage in harmful pricing tactics due to perverse incentives they would face if allowed to engage in fast lane pricing.

Uncertain Surplus Shift

Under prioritization, there would be a surplus shift in the Internet market; however, it is undeterminable how that redistribution would occur.


If ISPs develop Internet fast lanes, they would segment the graph above into two different markets. Each market would have its own demand curve. Figure 8 illustrates how the two markets might compare to the single lane market.

Take, as an example, an ISP charging content providers no price to access the slow lane and charging a price of P to place content in the fast lane.

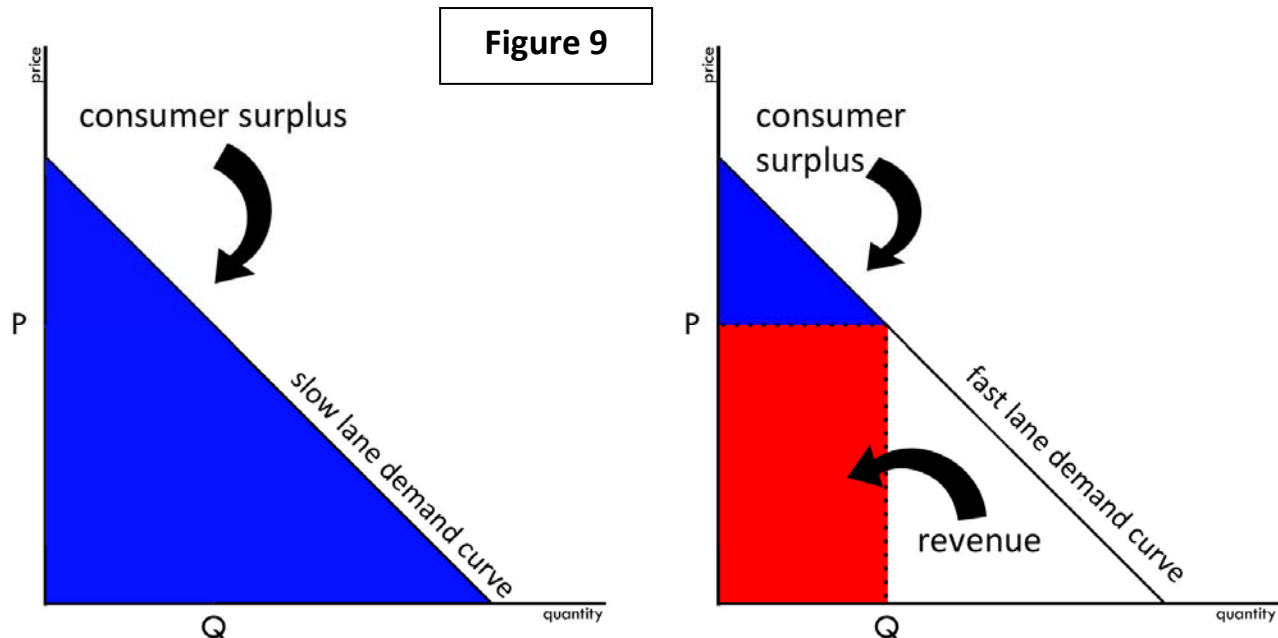
The content providers with a willingness-to-pay above the price will receive a higher level of service and therefore would have a higher willingness-to-pay for access to Internet subscribers. The content providers with a willingness-to-pay below price P to access the fast lane will be on the demand curve that shifts inward; they would receive a relatively lower level of service than content providers with a willingness-to-pay above the cutoff, and so their demand curve shifts inward.¹⁴⁵

Figure 9 below separates these markets into two demand curves, so we can more easily identify producer and consumer surplus in each market separately.

The demand curve on the left illustrates the demand for access to Internet subscribers over the slow lane and looks similar to a single stream Internet. The entire surplus in the market is captured by the Internet content providers. The graph on the right illustrates the consumer and producer surplus in the fast lane market, with ISPs capturing revenue.



The efficiency of the market is determined by the size of the total surplus generated by the single market compared to the sum of the surpluses created by the two (fast lane and slow lane) markets. The structure that generates the most surplus is most efficient. Unfortunately, there are no clear conclusions that can be drawn about the net effects of the creation of fast lane pricing without making additional assumptions about the markets.



Some obvious observations can be drawn about the impact of this switch on ISPs, content providers, and users. ISPs will clearly be better off if prioritization is allowed; they gained zero surplus from the single-stream market but would gain revenue by charging content providers for access to the fast lane.

Another clear observation: content providers with a willingness-to-pay for access to Internet subscribers below the price of access to the fast lane would be worse off. They would end up in the slow lane, and their surplus would decrease. Users would experience degraded connections to these websites during times of congestion. Some users may then stop using those sites—decreasing the surplus to the content providers in the slow lane. This may particularly be the case if the speed of information over the Internet is a “positional good”—i.e., a good that earns at least part of its value from being superior to comparable substitutes.¹⁴⁶ For example, users may prefer to use a search engine that retrieves search results faster. In that case, introducing a fast lane could reduce the willingness-to-pay of those in the slow lane even if their download speed has stayed the same or even increased.

As for the content providers in the fast lane, it is unclear how they would fare: they are worse off because they will be shifting some of their surplus to ISPs, but they would be better off relative to content providers in the slow lane. Further, if speed is a positional good, and content providers choose to switch from a websites on the slow lane to a faster one, some revenue that would have been gained by content providers in the slow lane will now be gained by content providers in the fast lane.

The effect on users is also unclear. Websites on the fast lane would presumably provide a higher quality experience, while websites on the slow lane may provide a worse experience than under the current single stream model. To the extent that the slow lane is slower or more interrupted than

current access speeds, it is difficult to predict the net impacts on consumers. In addition, if content providers leave the market in response to the creation of fast lanes, end-users would lose out. It is however possible that users may not actually perceive the reduction in quality of sites on the slow lane, producing none of these effects and making the outcome on users uncertain.¹⁴⁷

To ascertain the full impact of prioritization on the value of the Internet, it would be necessary to sum the reduction in surplus from slow lane content providers, the increase in surplus to ISPs, and the indeterminate change in surplus to fast lane content providers and users. Only if this sum were positive would prioritization increase the value of the Internet. It is unclear whether the winners in this scenario would be better off than the losers, and it is impossible *a priori* to determine how the market surplus would shift.

This sum is likely to be positive if the content providers in the fast lane have a significantly higher willingness-to-pay for access to Internet subscribers through a fast lane than through a single-stream Internet. It is also likely to be positive if the reduction in surplus is small to slow lane content providers. This would be the case if users of the websites in the slow lane are able to access those websites at the same quality and speed as they did in the single stream Internet, and if download speed is not a positional good. The sum is likely to be negative if the slow lane makes content providers significantly worse off, while the fast lane does not improve the user experience enough to generate a substantial increase in content providers' willingness-to-pay to access users.

Perverse Incentives

There is a reason to believe that if prioritization is allowed, ISPs will pursue it suboptimally. Because ISPs will be the market players controlling the prioritization pricing, they will be in control of the market surplus shift. As explained in Part Three, ISPs have incentives to maximize their own revenue, even if at the expense of the surplus of the Internet market as a whole. Not only do these incentives affect ISPs' execution of price discrimination, they also will affect ISPs' execution of prioritization pricing schemes.

Currently (and under a legal net neutrality regime), ISPs gain no revenue from content providers who access their subscribers, but the Internet generates significant value for content providers of all sorts, which ISPs cannot extract. Under a two-tier system, ISPs may be willing to enact pricing schemes that reduce the value of the Internet as a whole in order to gain a share of the benefits generated by content providers. In essence, ISPs currently receive none of a very large pie, but they would be willing to reduce the size of the pie in exchange for a slice.

Under most prioritization proposals, ISPs would offer slow lanes to content providers for free and charge a price for access to the fast lane. Because ISPs gain no revenue from content providers in the slow lane, they will have a strong incentive to maximize

their revenue from the fast lane. They could accomplish this by pricing fast lanes above marginal cost as explained above in order to price discriminate. Or, ISPs could attempt to push as many content providers into the fast lane as possible; and to incentivize more content providers to move into the fast lane, ISPs may compromise the quality of the slow lane.

ISPs will create prioritization policies based solely on their bottom lines, and will not consider the impact of this pricing on a sizeable fraction of the Internet market.

Because ISPs' incentives are not aligned with the market as a whole, prioritization schemes could increase ISPs' revenue, while lowering the total surplus value of the Internet.

The availability of prioritization pricing could incentivize ISPs to fail to improve the quality of the lower priced service—or even go so far as to decrease the quality of the lower priced good—so that consumers will upgrade to the premium service. ISPs may allow their lower (or standard) quality service to languish to encourage content providers to use the higher-priced fast lanes—so the ISP will gain more revenue from those who switchover. ISPs would have less incentive to upkeep the infrastructure needed to run their standard slow lines, and less incentive to build up that infrastructure.

Further, ISPs could actually degrade the performance of traffic running in the standard slow lane. There may be particular reason to think this scenario would occur given the state of technology. ISPs only have limited scope to speed up packets using prioritization.¹⁴⁸ The fast lane only provides speed boosts for content during times of congestion. During uncongested times, the fast lane would not be any faster than the slow lane. On the other hand, ISPs have the ability to slow traffic that runs over their network. To guarantee that fast lane service is faster than slow lane service, ISPs could degrade traffic in the slow lane.

A “best effort network” means that there is no guarantee that data is delivered or that a user receives a specific quality of service level or a certain priority.

In a best effort network, all users obtain the “best” bit rate and delivery time that the network can deliver, depending on the current traffic load.

If ISPs slow the service of the slow lane, content providers in the slow lane and users attempting to access those sites may be worse off, i.e. their surplus will decrease. Because those websites’ performance would be downgraded and even slower to load and access during times of congestion, users may choose not to access those sites and access competitors’ sites on the fast lane instead. If the content providers on the slow lane lose too many users, they may choose to exit the network.

Even if some content providers exit the market, and even if ISPs lose some subscribers as a result, ISPs’ revenues could still increase under this scenario. ISPs would not take into consideration if content providers in the slow lane lose surplus by dropping out of the market, because those players do not pay fees to ISPs. Because ISPs can generate revenue on both sides of the fast-lane market, but only from subscribers on the slow lane, they will have incentives to move content onto the fast lane even if that reduces overall surplus.¹⁴⁹

In this way, ISPs will not necessarily manage their prioritization schemes to maximize net surplus because they would likely not consider the impact of their prioritization policy on a sizeable fraction of the Internet market. This failure increases the probability that ISP-controlled prioritization schemes would lower the value of the Internet while increasing revenues for ISPs.

Effect on New Content

Supporters of prioritization argue that without it certain new and valuable high-bandwidth technologies will be unable to flourish, and the lack of these technologies will reduce the value of the Internet.¹⁵⁰ They fear these technologies (such as telemedicine) will never take off unless they can pay for high quality access to Internet subscribers because high bandwidth and high speed are necessary for those technologies.

Although this argument may have some validity, it is difficult to evaluate. It is hard to predict what types of technologies may develop in a prioritization market. There is a risk that net neutrality rules may prevent some beneficial technologies from developing—just as there is risk that prioritization would force some players out of the market that would have generated important innovations.

There are countless imaginable technologies that may someday require high levels of bandwidth, and there is no Internet policy regime that would not create some potential negative effects on these technologies. Of course, new content is only a portion of the total value of the Internet to society. The impact of fast lanes on ISPs, end-users, and existing content providers must also be considered when evaluating the ramifications of Internet fast lanes.

Moreover, the current “best effort network” has allowed a wide variety of technologies to proliferate. These technologies have been, and are being, created to run on the current Internet system. If the FCC believes that certain valuable technologies could not succeed on the best effort network, the FCC can choose to place these technologies into the “managed or specialized services” exception of the rule.¹⁵¹ The FCC, however, should narrowly construe that exception to only include technologies that cannot prosper over the best effort network.

Notably, ISPs are the most prominent supporters of this line of argument. The content companies attempting to develop these high bandwidth technologies have mixed opinions on the impact of net neutrality regulation on their business prospects.¹⁵² Further, many content providers and venture capitalists strongly endorse strict net neutrality rules, arguing that protecting the end-to-end architecture of the Internet encourages content entrepreneurship.¹⁵³ The fact that ISPs are the ones requesting fast lanes because they will help new companies but the venture capitalist that fund these companies appear not to want this service creates a quandary. It is possible that ISPs have a better sense of what conditions will generate new companies on the Internet than venture capitalists, but this may also be an example of ISPs attempting to acquire surplus from the market at the expense of content developers.

Free Prioritization

Some argue that ISPs should be allowed to prioritize and decide which sites to send to the slow lane and which to send to the fast lane so long as they do not charge content providers a price.¹⁵⁴ This practice could potentially be justified under the proposed rule.

Under the new rule, ISPs may engage in “reasonable network management” practices. The FCC will determine the full contours of what constitutes reasonable network management through case-by-case adjudication in the future. For example, the FCC could allow an ISP during times of congestion to prioritize packets from certain time-sensitive content (like streaming video or Voice-over Internet Protocol (“VoIP”)) for free if the ISP could prove that this prioritization is “reasonable network management.” The ISP would need to prove that practice was a “reasonable practice[]” to combat the “effects of congestion” or to “address quality-of-service concerns.”¹⁵⁵ Conceivably, ISPs could even prove free prioritization is “reasonable network management” during times of no congestion if they could prove doing so would improve quality of service.

In this way, if ISPs are truly concerned about congestion and quality of Internet access, they can convince the FCC to allow free prioritization schemes. Free prioritization implemented in this manner may not be detrimental from an economic perspective. It would not create the perverse incentives that paid prioritization generates, and would therefore not have the same reduction in total market surplus and problematic wealth transfer to ISPs.

For a variety of reasons, prioritization could run the risk of lowering the total surplus from the Internet market. If it turns into price discrimination in disguise, it will have all of the problems discussed above: additional surplus for ISPs will come at the expense of content providers, who are relatively more difficult to target for government support to overcome the positive externalities present in the Internet market. In addition, the pure surplus effects from breaking the Internet into multiple streams is uncertain, the surplus loss for content providers in the slow lane may offset any gains from content providers in the fast lane. Perhaps most importantly, ISPs will face perverse

incentives if they can generate revenue from the fast lane but not the slow lane. This misalignment of incentives could create a situation where ISPs can increase their revenue at the expense of the overall surplus from the market.

Conclusion

Toward a Smart Internet Policy

One of the difficulties in formulating and evaluating Internet policy is that the field is so new, but has come to be incredibly central to our economic and social lives. So while the stakes are very high, policymakers have relatively little experience or information on the likely effects of different policy approaches. There is no choice but to make the best decisions possible in the face of large uncertainties, knowing that the effects of our decisions will have profound social consequences.

This Report discusses some of the central questions that are raised by the FCC's proposed net neutrality rules and finds that, on balance, maintaining the current pricing scheme is likely to be economically justified. While opponents of net neutrality are correct that it may have some downsides—including decreased investment incentives for ISPs and potential impacts on technological development—the government has tools at its disposal to mitigate these downsides. Moreover, the benefits of net neutrality, especially maintaining investment incentives for the development of new content, are very high.

The goal of Internet policy should be the maximization of the surplus generated by the Internet market, understood as both physical infrastructure and the content travelling over that infrastructure. That surplus is vast, and only a portion of it is captured by private companies—Internet end-users derive significant surplus, and there are large positive externalities associated with Internet use. Protecting that surplus, and giving companies the right incentives to continue investing in increasing that surplus, should be the guiding principle of decisionmakers as they set Internet policy.

Ultimately, net neutrality represents something of a zero-sum game between ISPs and content providers. While moving away from net neutrality toward systems of price discrimination or prioritizations would benefit ISPs, much or all of that benefit would come at the expense of content providers. Because both ISPs and content providers produce positive externalities for society, benefiting one at the expense of the other is not sound policy.

Because of market failures that under-incentivize investment in the Internet, both infrastructure and content should be given government support to achieve optimal levels of investment. However, it is much easier to subsidize infrastructure than content—there is a long history of successful government support for infrastructure, while support for content in various forms has been riddled with controversy and difficulty. A policy that encourages content investment through a favorable pricing structure, while directly supporting infrastructure, then, is likely to be the best available option to achieve more efficient levels of investment.

As the government continues to try and give private actors the right incentives to continue investing in the Internet, it needs to retain flexibility to ensure that new technologies are appropriately promoted. The “reasonable network management” and “managed services” components of the rule, which are likely to be more fully explained in future rulemakings and adjudications, create this important flexibility. Maintaining a system that promotes broad investment, while ensuring that all important technologies are given room to grow is central to promoting the highest possible value of the Internet.

Notes

¹ Trans Atlantic Consumer Dialog, Resolution on Net Neutrality, at 2 (Mar. 2008), *available at* <http://www.publicknowledge.org/pdf/tacd-nn-resolution-200803.pdf>.

² Nicholas Economides, “Net Neutrality,” *Non-Discrimination and Digital Distribution of Content Through the Internet*, 4 I.S.J.L.P. 209, 211-12 (2008), *available at* http://www.stern.nyu.edu/networks/Economides_Net_Neutrality.pdf. *See* Part II.A for a more complete discussion of how end-users and content providers pay for Internet access.

³ Communications Act of 1934, Pub. L. No. 73-416, 48 Stat. 1064 (codified at 47 U.S.C. § 151 et seq.). The powers of the FCC were amended by the Telecommunications Act of 1996, Pub. L. No. 104-104, 110 Stat. 56, 71-72 (1996). The FCC is lead by five commissioners who are appointed by the President and confirmed by the Senate to serve five-year terms. Only three commissioners may be members of the same political party, and the President chooses one to serve as chairperson. 47 U.S.C. § 154.

⁴ 47 U.S.C. § 152. The FCC also has “ancillary authority to regulate interstate and foreign communications” even when not specifically delegated by statute, as well as broad authority to “conduct its proceedings in such manner as will best conduce to the proper dispatch of business and to the ends of justice,” 47 U.S.C. § 154(j); *see also* Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications, 23 F.C.C.R. 13,028, 13,044 (2008) [hereinafter Comcast Order]. The FCC has historically interpreted Internet policy as within the ambit of this ancillary authority. *See id.* at 13,036 (“Aside from section 230, we also find that exercising jurisdiction . . . is reasonably ancillary to our authority under . . . Section 1 of the Communications Act, section 201 of the Act, section 706 of the Act, and section 601(4) of the Act.”). Further, the Supreme Court has upheld the FCC’s ability to regulate in this area. *See, e.g., Nat’l Cable & Telecomms. v. Brand X Internet Services*, 545 U.S. 967, 996 (2005) (“[T]he Commission remains free to impose special regulatory duties on facilities-based ISPs under its Title I ancillary jurisdiction.”).

⁵ Comcast Order, *supra* note 4, at 13,045.

⁶ *Id.* In its Comcast Order, the FCC provided several reasons for this preference. First, because Internet and broadband policy are relatively new and thus “unsettled,” adjudication allows the agency the flexibility to shape policies according to new information, while still giving some guidance to consumers and industry

players. Second, network management practices are so “complex and variegated” that general rules may be difficult, or even undesirable. Finally, adjudication adheres more closely to “congressional directives and Commission precedents.” *Id.* at 13,045-46.

⁷ FCC, Appropriate Framework for Broadband Access to the Internet over Wireline Facilities: Inquiry Concerning High-Speed Access to the Internet over Cable and Other Facilities, Policy Statement, 20 F.C.C.R. 14,986, 14,988 (2005) [hereinafter Internet Policy Statement]; *see also* FCC, Broadband Industry Practices, Notice of Inquiry, 22 F.C.C.R. 7894, 7896 (2007) [hereinafter Notice of Inquiry] (“The Commission, under Title I of the Communications Act, has the ability to adopt and enforce the net neutrality principles it announced in the Internet Policy Statement.”).

⁸ Internet Policy Statement, *supra* note 7, at 14,988. In 2005, Vonage complained to the FCC that Madison River Communications (an ISP) blocked its customers’ use of Vonage’s Voice-over Internet Protocol (“VoIP”), an application that converts voice into digital signals to travel over the Internet. Madison River and the FCC then entered into a Consent Decree, fining Madison River and prohibiting it from engaging in this type of discrimination. *Madison River Commc’ns LLC and Affiliated Cos.*, 20 F.C.C.R. 4295 (2005). Although this was a consent decree and the FCC did not explain its position, it nevertheless represents a general commitment to net neutrality principles.

⁹ Notice of Inquiry, *supra* note 7, at 7,905 (concurring statement of Comm’r Jonathan S. Adelstein).

¹⁰ Comcast Order, *supra* note 4, at 13,028.

¹¹ *Id.* at 13,050–51.

¹² *Id.* at 13,044.

¹³ *Id.* at 13,052 (quoting Internet Policy Statement, *supra* note 7, at 14,988) (internal punctuation omitted).

¹⁴ *Comcast Corp. v. FCC*, No. 08-1291 (D.C. Cir., filed Sept. 4, 2008). Oral argument is scheduled for January 8, 2010. The case considers whether the FCC has authority to require verification that Comcast had discontinued its practice of blocking applications, and require Comcast to disclose any other practice that prevents customer access of Internet applications; and whether the FCC permissibly considered Comcast’s actions in an adjudication rather than a rulemaking.

¹⁵ Preserving the Open Internet, Broadband Industry Practices, 74 Fed. Reg. 62,638, 62,638 (proposed Nov. 30, 2009) (to be codified at 47 C.F.R. pt. 8) [hereinafter Broadband Industry Practices].

¹⁶ *Id.* at 62,644.

¹⁷ *Id.* at 62,649.

¹⁸ *Id.* at 62,654.

¹⁹ *Id.* at 62,644.

²⁰ *Id.* at 62,646, 62,651–52.

²¹ *Id.* at 62,644–46.

²² *Id.* at 62,646–47.

²³ *Id.* at 62,647–48.

²⁴ *Id.* at 62,661.

²⁵ *Id.* at 62,646.

²⁶ *Id.*

²⁷ *Id.* at 62,661.

²⁸ *Id.* at 62,650.

²⁹ *Id.*

³⁰ *Id.* (“[O]pen Internet principles apply only to lawful transfers of content. They do not, for example, apply to activities such as the unlawful distribution of copyrighted works . . .”).

³¹ *Id.* at 62,650–51 (“[W]e believe that additional flexibility to engage in reasonable network management provides network operators with an important tool to experiment and innovate as user needs change.”).

³² *Id.* at 62,651 (proposing rules that clarify that broadband providers maintain a superseding obligation to law enforcement, public safety, and national security concerns).

³³ A full analysis of the impact of net neutrality on the value of the Internet would require information on the value of the Internet under net neutrality and the value of the Internet under a price discrimination regime. The Internet currently operates under a *de facto* net neutrality regime, and its costs and benefits could be valued. However, estimating the value of the Internet under price discrimination would require a series of heroic assumptions about how ISPs and content providers would respond to price discrimination. A specific estimate of the value of the Internet under price discrimination would be speculative. It would therefore be difficult to compare the benefit and cost ratios of both regimes.

³⁴ Brett M. Frischmann & Barbara van Schewick, *Net Neutrality and the Economics of an Information Superhighway: A Reply to Professor Yoo*, 47 JURIMETRICS J. 383, 424–27 (2007) (arguing that net neutrality fosters innovation “critical for economic growth,” and that the nondiscriminatory aspect of net neutrality “may create enormous social value”).

³⁵ When an Internet subscriber requests a webpage from a browser, her computer sends a signal to the location where that website is stored. All the information necessary for the browser to display the website is broken up into tiny pieces of information called “packets,” which are sent across the Internet from the website’s location to the Internet subscriber’s computer. The ISP’s router acts as a “gate” from the larger Internet network to its private last mile network. At this gate, the ISP can determine the destination of the packet within its private network. The router then sends the packet of information over the ISP’s private network to the computer of the Internet subscriber requesting the website. Finally, the user’s web browser reassembles the packets to form a webpage.

³⁶ For a more detailed explanation of the structure of Internet, see PRESTON GALA, *HOW THE INTERNET WORKS* (1998).

³⁷ Robert Hahn & Scott Wallsten, *The Economics of Net Neutrality*, 3 ECONOMISTS’ VOICE 1 (2006), available at <http://www.bepress.com/ev/vol3/iss6/art8>.

³⁸ S. J. Liebowitz & Stephen E. Margolis, *Network Externality: An Uncommon Tragedy*, 8 J. ECON. PERSP. 133, 135–136 (1994) (explaining the economics of the network externality).

³⁹ David D. Kirkpatrick, *As Instant Messaging Comes of Age, AOL Fears F.C.C. Rule Holds It Back*, N.Y. TIMES, May 26, 2003, available at <http://www.nytimes.com/2003/05/26/business/as-instant-messaging-comes-of-age-aol-fears-fcc-rule-holds-it-back.html> (describing how the FCC, “[a]s a condition of approving AOL’s acquisition of Time Warner,” made it more difficult for AOL to offer video chat capabilities to users if AOL did not link its instant messaging with competitors).

⁴⁰ See, e.g., Andrew Caplin & John Leahy, *Miracle on Sixth Avenue: Information Externalities and Search*, 108 ECON. J. 60, 61, 71–72 (1998).

⁴¹ See, e.g., J. Vickers, *Regulation, Competition, and the Structure of Prices*, 13 OXFORD. REV. ECON. POL’Y 15, 15–19 (1997). From an economic perspective, price regulation may be preferred to “quantity regulation” (output caps) or “rate of return regulation.” See Gary Biglaiser & Michael Riordan, *Dynamics of Price Regulation*, 31 RAND J. ECON. 744 (2000), for further discussion.

⁴² See, e.g., GREGORY MANKIW, *PRINCIPLES OF MICROECONOMICS* 114–23 (5th ed. 2008).

⁴³ See Martin Neil Baily, Robert J. Gordon & Timothy F. Bresnahan, *Competition, Regulation, and Efficiency in Service Industries*, 1993-2 BROOKINGS PAPERS ON ECONOMIC ACTIVITY: MICROECONOMICS 71, for one empirical assessment and reference others.

⁴⁴ See, e.g., Liebowitz & Margolis, *supra* note 38, at 137.

⁴⁵ See John Musacchio, Galina Schwartz & Jean Walrand, *A Two-Sided Market Analysis of Provider Investment Incentives With an Application to the Net-Neutrality Issue*, 8 REV. NETWORK ECON. 22, 34–36 (2009). Nicholas Economides contends that it is incorrect to conceptualize the Internet as a two-sided market with ISPs in the middle and content providers and end-users on each side. He argues instead that one side of the

market consists of Internet subscribers who pay ISPs for access; and on the other side are Internet backbone companies who enter into transit agreements with ISPs to carry traffic to and from the ISPs' private networks. Content providers are not part of this market. If ISPs believe it would be profitable to increase investment in their network infrastructure, they are free to raise revenue by increasing the prices they charge subscribers or the prices they charge backbone companies and use that revenue to invest in infrastructure (or return it to their shareholders, or for whatever purpose they deem fit). If ISPs increase the prices they charge backbone companies, those companies could then increase the prices they charge to other companies that use their lines, including content providers. In this way, ISPs should use the *existing* market structure to raise revenue as opposed reaching outside the market to content providers for revenue. Economides explains that if ISPs begin charging content providers directly they will bypass the market structure that exists and take advantage of their exclusive access to end-users and transfer revenue to themselves from other efficient markets. See Nicholas Economides & J. Tag, *Net Neutrality on the Internet: A Two-sided Market Analysis* 6 (NYU Stern School of Business 2009), available at http://www.stern.nyu.edu/networks/Economides_Tag_Net_Neutrality.pdf. Even if it is correct to think of the Internet market as two-sided with end-users on one side and content providers on the other, it is not clear that in general equilibrium, allowing network owners to extract surplus from content providers would increase aggregate surplus. By increasing the rate of return on investment in the network, it would likely incentive expansion of the network, increasing surplus for network owners and content providers. However, the capital used to build the network would come from elsewhere in the economy, where presumably it would be making market rates of return—the shift in capital from those markets to the broadband market does not necessarily represent an aggregate increase in surplus. Because the Internet does face pervasive positive externalities, a shift in capital may be warranted, but, as discussed elsewhere in this report, price discrimination is a blunt tool to achieve this goal, because funds are extracted from content providers (who also generate positive externalities), and direct government support can be better targeted to maximize the surplus benefits of investment.

⁴⁶ Economides, *Net Neutrality*, *supra* note 45, at 211 (noting the claim that consumers, content providers, or applications providers use the Internet for free is certainly incorrect).

⁴⁷ Paolo Buccirossi, Laura Ferrari Bravo & Paolo Sicilani, *Competition in the Internet Backbone Market*, 2 *WORLD COMPETITION* 235, 243-345 (2005).

⁴⁸ See Tony Chan, *Does Google Pay Its Fair Share for Bandwidth?*, COMMSDAY.COM, Dec. 10, 2008, <http://www.commsday.com/node/295>.

⁴⁹ Frischmann & van Schewick, *supra* note 34, at 415.

⁵⁰ Hahn & Wallsten, *supra* note 39, at 37.

⁵¹ See Patricia O'Connell, *At SBC, It's All About "Scale and Scope,"* BUSINESS WEEK, Nov. 7, 2005, http://www.businessweek.com/@n34h*IUQu7KtOwgA/magazine/content/05_45/b3958092.htm (Edward Whitacre, CEO of AT&T, responding to the question, "How concerned are you about Internet upstarts like Google, MSN, Vonage, and others?").

⁵² Burkhard Stiller & Peter Reichl, *Pricing and Cost Recovery for Internet Services: Practical Review, Classification, and Application of Relevant Models*, 3 *NETNOMICS* 149, 166-168 (2001). This is a somewhat simplified example: the ISPs may not pay based on quantity of traffic that traverses their network, but on the value of traffic which includes both quantity and distance traveled. *Id.* at 166.

⁵³ One terabyte is 1 trillion bytes or 1,000 gigabytes; it can hold data the equivalent of approximately 260,000 average length songs.

⁵⁴ At regular intervals, backbone servers note how much traffic they sent and received to other families of servers. If they sent more data than they received, they remit a payment to the other network. If they received more data than they transmitted to a particular backbone server or ISP, the other network makes a payment to them. See Stiller & Reichl, *supra* note 52, at 150-151.

⁵⁵ M. CHRIS RILEY & BEN SCOTT, FREE PRESS, DEEP PACKET INSPECTION: THE END OF THE INTERNET AS WE KNOW IT? 3 (March 2009), http://www.freepress.net/files/Deep_Packet_Inspection_The_End_of_the_Internet_As_We_Know_It.pdf.

⁵⁶ *Id.*

57 GEORGE OU, DIGITAL SOCIETY, UNDERSTANDING DEEP PACKET INSPECTION (DPI) TECHNOLOGY 5-6 (2009), <http://www.digitalsociety.org/files/gou/DPI-Final-10-23-09.pdf> (describing Denial of Service attacks as a “classic example” of a situation where network operators can use DPI to minimize false positives and negatives).

58 *Id.* at 6 (referencing an October 2009 announcement by Comcast).

59 *Id.* at 7 (referencing the Communications Assistance for Law Enforcement Act of 1994).

60 *Id.*

61 *Id.*

62 S. DEREK TURNER, FREE PRESS, DISMANTLING DIGITAL REGULATION: TOWARD A NATIONAL BROADBAND STRATEGY 75 (2009), http://www.freepress.net/files/Dismantling_Digital_Deregulation.pdf (noting that the current broadband market is a “duopoly” in which phone and cable companies control “97 percent of the fixed-line residential broadband market”).

63 Riley & Scott, *supra* note 55, at 3.

64 The FCC itself realized this possibility when it noted that, “With ‘deep packet inspection,’ . . . [a] broadband Internet access service provider can . . . favor certain parties by providing access to information cached at the provider’s facility, allowing consumers quicker access to Web sites . . .” Broadband Industry Practices, *supra* note 15, at 62,641.

65 As explained in the Introduction, a Peer-to-Peer Application allows participants to make a portion of their resources directly available to other participants, without the need for central coordination. *See supra* note 11 and accompanying text.

66 Comcast Order, *supra* note 4, at 13,051 (“When Comcast judges that there are too many peer-to-peer uploads in a given area, Comcast’s equipment terminates some of these connections by sending RST packets”).

67 *Id.* at 13,028 (requiring Comcast to, *inter alia*, terminate its current DPI practices).

68 See TURNER, *supra* note 62, at 75, for more discussion on Comcast’s practices.

69 *Id.*

70 *Id.* In October 2009, after the company tested DPI monitoring in Kansas and Arkansas, it announced a decision against network-wide implementation. *See Cox Communications, Congestion Management FAQs*, Dec. 29, 2009, <http://www.cox.com/policy/congestionmanagement>.

71 See TURNER, *supra* note 62, at 76 (noting that Primus has imitated Cox’s system in Canada).

72 Ou, *supra* note 57, at 4.

73 Grant Gross, *Deep Packet Inspection Could Be Outlawed in U.S.*, TECH WORLD, Apr. 24, 2009, <http://news.techworld.com/security/114856/deep-packet-inspection-could-be-outlawed-in-us>.

74 Lars Stole, *Price Discrimination in Competitive Environments*, in 3 HANDBOOK OF INDUS. ORG. 1, 1-3 (Mark Armstrong & Robert H. Porter eds., 2007).

75 N. GREGORY MANKIW, PRINCIPLES OF MACROECONOMICS 403 (5th ed. 2008). In contrast, in competitive markets, arbitrage leads to equal prices for the same goods.

76 Joshua S. Gans & Stephen P. King, *Perfect Price Discrimination with Costless Arbitrage* 16-18 (MELBOURNE BUSINESS SCHOOL WORKING PAPER No. 15 2005).

77 Marius Schwartz, *Third-Degree Price Discrimination and Output: Generalizing a Welfare Result*, 80 AM. ECON. REV. 1259 (1990).

78 The analog is “producer surplus,” which is generated for producers when they are able to charge a higher market price than their minimum “willingness-to-supply” level.

79 In economic theory, wealth distribution should not affect the efficiency of the market. Economic efficiency is a positive (or fact-based question), while the optimal distribution of benefits from a market is a normative (or value-based question). *But see* JURGEN HABERMAS (WILLIAM REHG, TRANS.), BETWEEN FACTS AND NORMS (2006) (arguing the distinction between positive and normative questions may not be clear cut). It may be the case that the fairest distribution of benefits is one in which each individual in the market gets an

equal share of the rents; or, one could argue, that the rents generated by a market should be distributed in proportion to the effort of each individual involved in the market, or perhaps in proportion to ability. Reasonable people may disagree about the correct way to distribute the benefits between groups.

Economists typically focus their attention on market efficiency because it has a simple and straightforward definition (maximizing the economic rents produced by a market), while the distribution of those rents is often noted but is typically not the focus of analysis. Economists can use sophisticated tools to measure the rents generated by a market and can develop precise estimates of market efficiency; however, the tools to determine optimal distribution of benefits from a market are not as sophisticated.

⁸⁰ Christing Zhen-Wei Qiang, Carlo Rossotto, & Kaoru Kimura, *Economic Impacts of Broadband*, in INFORMATION AND COMMUNICATIONS FOR DEVELOPMENT 2009: EXTENDING REACH AND INCREASING IMPACT, WORLD BANK 53 (2009). This is a linear estimate from a standard Ordinary Least Squares (OLS) regression.

⁸¹ It is possible that broadband penetration is correlated with some other variable (such as average education levels or institution quality) that causes economic growth. If this variable is not observed, then it may appear that broadband penetration causes growth when in fact the third variable is generating both economic growth and increased broadband penetration. For example, there are a variety of differences in legal systems, cultural norms, and institutions between higher broadband countries and low broadband countries—all could be confounding variables, and these differences are extremely difficult to control for statistically. Countries with high levels of broadband penetration also tend to have higher education levels, initial incomes, and many other hard-to-measure differences from countries that have lower levels of broadband uptake. Even with this caveat, information on the Internet is clearly extremely valuable, and improving access to that information will certainly have some cognizable economic benefits.

⁸² To calculate this number exactly requires determining each Internet subscriber's willingness-to-pay for Internet access and comparing their maximum willingness-to-pay to the actual cost of subscription. The difference between these numbers summed across all Internet subscribers is the total surplus. Because finding the willingness-to-pay of each subscriber is extremely difficult, we are left to estimate the surplus generated by the Internet using imprecise measures.

⁸³ See, e.g., Yochai Benkler, *Next Generation Connectivity: A Review of Broadband Internet Transitions and Policy from Around the World* 10, 13 (Berkman Center at Harvard University, 2009), available at http://www.fcc.gov/stage/pdf/Berkman_Center_Broadband_Study_13Oct09.pdf [hereinafter Berkman Report] (arguing that broadband access in the United States is relatively poor compared to other wealthy countries). For criticism of the report, see Seth L. Cooper, *Free State Foundation, The Faulty Berkman Report: The Fallacy of Overlooking Secondary Consequences*, PERSPECTIVES FROM FSF SCHOLARS, Nov. 3, 2009, available at http://www.freestatefoundation.org/images/The_Faulty_Berkman_Report.pdf (arguing that attempts to compare broadband policy across countries are flawed).

⁸⁴ Berkman Report, *supra* note 83, at 13. The Berkman Report makes its findings after controlling for a host of differences. However, regardless of whether or not other differences are controlled for, the United States remains in the middle of the pack. *Id.* at 26-39.

There is considerable debate about the quality of the OECD data on which this analysis is based and the statistical techniques used in the Berkman Report. See, e.g., Comments from AT&T Inc, et al., to FCC, on National Broadband Plan Public Notice # 13, on Berkman Center Report (Nov. 16, 2009) (in public docket for GN Docket Nos. 09-47, 09-51, 09-137) (disputing Report's data), available at, <http://fjallfoss.fcc.gov/ecfs/document/view;jsessionid=LhhLCzjgfvGx4dkPnxwhVyyqkCn1z1wf7mKWm9pypkX5P79PswqCr!1096618407!1219827282?id=7020348449>.

⁸⁵ Berkman Report, *supra* note 83, at 13-14. As mentioned, some have taken issue with the Berkman Center Report's statistical techniques and the OECD data on which its analysis relies. See, e.g., George Ou, Digital Society, *Flawed Data in Berkman Broadband Study*, <http://www.digitalsociety.org/2009/10/flawed-data-in-berkman-broadband-study> (Oct. 19, 2009).

⁸⁶ This value represents the market capitalizations for Google, eBay, Yahoo, Amazon and IAC (owner of Ask.com and other sites). See <http://finance.yahoo.com> (last visited Dec. 2, 2009). Facebook likely belongs on this list, but it is privately held. Many other large companies that derive less than 100% of their market

capitalization from Internet content provision were excluded. “Market capitalization” refers to the value of a public corporation’s price per share multiplied by the number of outstanding shares.

⁸⁷ Alex Goldman, *Top 23 U.S. ISPs by Subscriber: Q3 2008*, INTERNET.COM, Dec. 2, 2008, <http://www.isp-planet.com/research/rankings/usa.html>.

⁸⁸ *Id.*

⁸⁹ Letter from Karen Rehfeld, Director, State of Alaska Office of Management and Budget, to Broadband Technology Opportunities Program (Apr. 9, 2009), *available at* http://gov.state.ak.us/omb/10_omb/budget/Economic_Stimulus/BTOP%20Broadband%20comments%204-9-09.pdf (noting that “[w]hen it comes to access to broadband, Alaska residents are the most ‘unserved’ and ‘underserved’ population in the United States”).

⁹⁰ TURNER, *supra* note 62, at 8 (noting that the current broadband market is a “duopoly” in which phone and cable companies control “97 percent of the fixed-line residential broadband market”).

⁹¹ *Id.* at 59 (finding that “barriers to entry in the telecom market are truly insurmountable”).

⁹² *See id.* at 55 & n.162.

⁹³ Because the Internet service provision market has a high barrier to entry, ISPs may be able to charge prices above equilibrium without attracting additional competition. While there are a large number of ISPs, in local markets there are often oligopoly conditions. Some argue that a net neutrality regime protects against abuse of market power by ISPs. *See, e.g.*, Berkman Report, *supra* note 82, at 16.; BEN SCOTT, MARK COOPER & JEANNINE KENNEY, WHY CONSUMERS DEMAND INTERNET FREEDOM: NET NEUTRALITY: FACT VS. FICTION 16 (2006), http://www.freepress.net/files/nn_fact_v_fiction_final.pdf. ISPs respond that FCC policy is not the forum to deal with possible antitrust issues. If the exercise of market power in the Internet service market causes a reduction in the value of the Internet, ISPs argue it should be dealt with through antitrust litigation or through a body like the Federal Trade Commission. This issue represents an important component of the discussion over net neutrality, but is not a focus of this Report.

⁹⁴ The vast majority of households are served by two ISPs and literally millions of websites. In any given Internet content industry (search engines, news sources, vendors) there are significantly more than two competitors. *See* TURNER, *supra* note 62, at 8 (noting that the current broadband market is a “duopoly” in which phone and cable companies control “97 percent of the fixed-line residential broadband market”).

⁹⁵ The concept of willingness-to-pay is more fully discussed in Part II.C.

⁹⁶ This is obviously a stylized example. In the case of a popular site like Wikipedia, potential losses in subscribers would force ISPs to engage in specific negotiations, and would give the content provider significant bargaining leverage. But, because the parties do not face all of the losses if a bargain is not reached (due to positive externalities) there is a strong possibility that social beneficial transactions will not occur.

⁹⁷ *See* Hahn & Wallsten, *supra* note 37, at 5.

⁹⁸ *See* George S. Ford, Thomas M. Koutsy & Lawrence J. Spiwak, *The Efficiency Risk of Net Neutrality Rules* 16 (Phoenix Ctr. Policy Bulletin, Working Paper No. 16, 2006), *available at* <http://www.phoenix-center.org/PolicyBulletin/PCPB16Final.pdf>.

⁹⁹ This is not always the case; climate change is a notable example. *See, e.g.*, Comments from Institute for Policy Integrity and Environmental Defense Fund, to Carol M. Browner, Assistant to the President, Office of Energy and Climate Change Policy on Federal Interagency Review on the Social Cost of Carbon 6-7 (Sept. 11, 2009) (discussing uncertainty and costs of climate change).

¹⁰⁰ It could be argued that ISPs should be able to charge content providers their marginal costs for providing access to their last mile, reflected in the congestion caused by content flowing over their networks. However, ISPs can better recoup those marginal costs from end-users—there is no reason to charge content providers when it is end-users who have ultimate control over what content travels over the ISPs’ networks. *See* Frischmann & Schewick, *supra* note 34, at 383–428; Christopher S. Yoo, *Network Neutrality and the Economics of Congestion*, 95 GEO. L.J. 1847, 1902 (2006).

¹⁰¹ Yoo, *supra* note 100, at 1884.

¹⁰² *See* Brett Frischmann, INFRASTRUCTURE: THE SOCIAL VALUE OF SHARED GOODS 11-28 (Yale University Press, forthcoming 2010), for a full analysis of the economics of infrastructure.

- ¹⁰³ Some nonrival goods can be subject to congestion: after a certain point they become rival as additional users consume the good. A highway that is one car short of a traffic jam is an example. The Internet does face congestion issues and is therefore not perfectly nonrival. Some have argued that price discrimination helps account for congestion over the Internet, but that argument has been strongly refuted. See generally Frischmann & van Schewich, *supra* note 34; Yoo, *supra* note 100.
- ¹⁰⁴ Richard S. Whitt, *Evolving Broadband Policy: Taking Adaptive Stances to Foster Optimal Internet Platforms*, 17 COMMLAW CONCEPTUS 17-18 (2009).
- ¹⁰⁵ “Fixed costs” are those that do not vary with output. For example a restaurant’s rent is a fixed cost because it does not change as the quantity of diners served changes. In the Internet market, the cost of setting up a network is considered a fixed cost. See *id.* for further explanation.
- ¹⁰⁶ Guy Michaels, *The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System*, 90 REV. ECON. & STAT. 683, 683-85 (2008).
- ¹⁰⁷ H. S. Person, *The Rural Electrification Administration in Perspective*, 24 AGRIC. HISTORY 70, 70 (1950).
- ¹⁰⁸ Edward M. Gramlich, *Infrastructure Investment: a Review Essay*, 32 J. ECON. LIT. 1176, 1185-89 (1994).
- ¹⁰⁹ RICHARD F. WEINGROFF, FED. HIGHWAY ADMIN., U.S. DEP’T OF TRANSP., ORIGINAL INTENT: PURPOSE OF THE INTERSTATE SYSTEM 1954-1956, available at <http://www.fhwa.dot.gov/infrastructure/originalintent.cfm> (last visited Jan. 5, 2010).
- ¹¹⁰ WENDELL COX & JEAN LOVE, AMERICAN HIGHWAY USERS ALLIANCE, 40 YEARS OF THE US INTERSTATE HIGHWAY SYSTEM: AN ANALYSIS: THE BEST INVESTMENT A NATION EVER MADE 2 (1996), <http://www.publicpurpose.com/freeway1.htm> (noting that the Dwight D. Eisenhower System of Interstate and Defense Highways has “returned more than \$6 in economic productivity for each \$1 it cost”).
- ¹¹¹ Amitabh Chandra & Eric Thomspon, *Does Public Infrastructure Affect Economic Activity? Evidence From The Rural Interstate Highway System*, 30 REGIONAL SCI. & URB. ECON. 457, 481 (2000).
- ¹¹² COX & LOVE, *supra* note 110, at 4.
- ¹¹³ *Id.* at 3.
- ¹¹⁴ See New Deal Network, *TVA: Electricity for All*, <http://newdeal.feri.org/tva/tva10.htm> (last visited Jan. 5, 2010) (noting that the Electric Home and Farm Authority, established by the TVA, helped farmers to purchase “major electric applications”).
- ¹¹⁵ Person, *supra* note 107, at 75.
- ¹¹⁶ New Deal Network, *supra* note 114 (“By 1939 rural households with electricity had risen to 25 percent.”).
- ¹¹⁷ Communications Act of 1934, Pub. L. No. 73-416, 48 Stat. 1064.
- ¹¹⁸ Lynne Holt & Mark A. Jamison, *Re-evaluating FCC Policies Concerning the Lifeline & Linkup Programs*, 5 J. ON TELECOMM. & HIGH TECH. L. 393, 393 (2007).
- ¹¹⁹ 47 U.S.C. § 254.
- ¹²⁰ FCC, *Universal Service*, http://www.fcc.gov/wcb/tapd/universal_service (last visited Jan. 5, 2009). Unlike the other programs, Lifeline and Link-up, which provide direct subsidies to low-income users, currently only subsidize traditional telephone service, though pending bipartisan Senate legislation would include support for broadband Internet access. Andrew Feinberg, *Rockefeller, Hutchison Would Expand Lifeline Program to Broadband*, BROADBANDBREAKFAST.COM, Dec. 11, 2009, <http://broadbandbreakfast.com/2009/12/rockefeller-hutchison-would-expand-lifeline-program-to-broadband>.
- ¹²¹ Jonathan S. Adelstein, Preface, 13 CATH. U. J. COMM. L. & POL’Y 1, 2 (2004) (former FCC Commissioner).
- ¹²² Press Release, White House, Vice President Biden Kicks Off \$7.2 Billion Recovery Act Broadband Program (Dec. 17, 2009), available at <http://www.whitehouse.gov/the-press-office/vice-president-biden-kicks-72-billion-recovery-act-broadband-program> [hereinafter White House Press Release].
- ¹²³ See Nat’l Telecomm. & Info. Admin., U.S. Dep’t of Commerce, *Broadband Technology Opportunities Program*, <http://www.ntia.doc.gov/broadbandgrants/> (last visited Jan. 5, 2010).

-
- ¹²⁴ Nat'l Telecomm. & Info. Admin., U.S. Dep't of Commerce, *Broadband Technology Opportunities Program Notice of Funds Availability Fact Sheet*, http://www.ntia.doc.gov/broadbandgrants/BTOP_NOFAfactsheet_090702.pdf (last visited Jan. 5, 2010).
- ¹²⁵ See Nat'l Telecomm. & Info. Admin., U.S. Dep't of Commerce, *Broadband Initiatives Program*, <http://www.broadbandusa.gov/files/BroadbandInitiativesProgramFactSheet8-12.pdf> (last visited Jan. 5, 2010) (detailing Broadband Initiatives Program). There is some overlap between these two provisions, as some initial recipients of BTOP funds are in rural areas. See White House Press Release, *supra* note 122.
- ¹²⁶ See Celia Kang, *FCC Issues Proposals to Meet National Broadband Plan*, WASH. POST, Dec. 17, 2009, available at <http://www.washingtonpost.com/wp-dyn/content/article/2009/12/16/AR2009121603916.html>.
- ¹²⁷ *Id.* Among the proposals is one that would redirect Universal Service Fund money in order to provide access to rural areas—exactly what the other two stimulus programs do. *Id.*
- ¹²⁸ White House Press Release, *supra* note 122.
- ¹²⁹ National Economic Counsel, RECOVERY ACT INVESTMENTS IN BROADBAND: LEVERAGING FEDERAL DOLLARS TO CREATE JOBS AND CONNECT AMERICA 11 (Dec. 2009), available at <http://www.whitehouse.gov/sites/default/files/20091217-recovery-act-investments-broadband.pdf>.
- ¹³⁰ Depending on how the government implements support of content, a program to subsidize Internet content may create potential First Amendment issues. See, e.g., Paul Farhi, *PBS Blesses Old Religious Shows, But Bans the New*, WASH. POST, June 17, 2009, <http://www.washingtonpost.com/wp-dyn/content/article/2009/06/16/AR2009061603201.html> (discussing debate over First Amendment and religious programming on PBS).
- ¹³¹ *Thematic Window: The Corporation for Public Broadcasting*, PUBLIC BROAD. SERV., <http://www.pbs.org/johngardner/chapters/4d.html> (last visited Jan. 5, 2010) ("Since the 1940's, the federal government has supported public stations because of their educational value.").
- ¹³² CORP. FOR PUBLIC BROAD., PUBLIC BROADCASTING REVENUE: FISCAL YEAR 2008 Table 1 (2009), <http://www.cpb.org/stations/reports/revenue/2008PublicBroadcastingRevenue.pdf> (categorizing funding by source).
- ¹³³ Steve Rendall & Peter Hart, *Time to Unplug the CPB: Replace Corrupt Board with Independent Trust*, FAIRNESS AND ACCURACY IN REPORTING, Sept.–Oct. 2005, <http://www.fair.org/index.php?page=2671> (describing, *inter alia*, Corporation for Public Broadcasting Chair Kenneth Tomlinson's "charges about the liberal bias of public broadcasting").
- ¹³⁴ Cliff Kincaid, *The Case for De-Funding Public Broadcasting*, ACCURACY IN MEDIA, June 21, 2005, <http://www.aim.org/special-report/the-case-for-de-funding-public-broadcasting> (arguing that "[t]he rationale for taxpayer funding of public broadcasting doesn't make sense . . . with more media choices than ever before").
- ¹³⁵ William H. Honan, *Book Discloses That Reagan Planned To Kill National Endowment for Arts*, N.Y. TIMES, May 15, 1988, at 155, available at <http://www.nytimes.com/1988/05/15/arts/book-discloses-that-reagan-planned-to-kill-national-endowment-for-arts.html?pagewanted=1>.
- ¹³⁶ *National Endowment for the Arts v. Finley*, 524 U.S. 569, 578 (1998) (describing how the NEA agreed to pay respondents an amount equal to the vetoed grants, damages, and attorney's fees).
- ¹³⁷ M. CHRIS RILEY, FREE PRESS, & ROBB TOPOLKI, NEW AMERICAN FOUNDATION, THE HIDDEN HARMS OF APPLICATION BIAS 1 (Nov. 2009), http://www.freepress.net/files/The_Hidden_Harms_of_Application_Bias.pdf (arguing that prioritization "poses hidden harms for the Internet that substantially outweigh its uncertain benefits").
- ¹³⁸ See generally George Ou, Digital Society, *Debunking the Myth That Prioritized Networks Are Harmful*, (Nov. 12, 2009), <http://www.digitalsociety.org/2009/11/debunking-the-myth-that-prioritized-networks-are-harmful> (arguing that prioritized networks are an efficient and fair way to manage the network).
- ¹³⁹ Broadband Industry Practices, *supra* note 7, at 62,646.

¹⁴⁰ See Vishesh Kumar & Christopher Rhoads, *Google Wants Its Own Fast Track on the Web*, WALL STREET J., Dec. 15, 2008, at A1, *available at* <http://online.wsj.com/article/SB122929270127905065.html> (quoting Ben Scott of Free Press).

¹⁴¹ TURNER, *supra* note 62, at 77 (arguing that fast lane pricing “raises the costs of entry, increases costs for consumers, and turns the Internet into a form of pay-for-play media”).

¹⁴² Scott et al, *supra* note 93 (“Network discrimination through a ‘tiered Internet’ will severely curtail consumer choice, giving consumer control over the Internet to the network owners”).

¹⁴³ See Robert E. Litan, *Catching the Web in a Net of Neutrality*. WASH. POST, May 2, 2006, *available at* <http://www.washingtonpost.com/wp-dyn/content/article/2006/05/01/AR2006050101061.html>. (discussing who should pay for fast lanes).

¹⁴⁴ Robert B. Ekelund, Jr., *Price Discrimination and Product Differentiation in Economic Theory: An Early Analysis*, 84 Q. J. OF ECON. 1259, 1276 (1970).

¹⁴⁵ The same graph could be drawn with no reduction in willingness-to-pay for content providers on the slow lane. This could happen if, for example, users’ experience of content on the slow lane was not degraded—this would happen if speed was not a positional good for Internet content. Even without this reduction, the value of the Internet may still decrease as a whole because of perverse incentives generated by prioritization schemes. These perverse incentives are discussed in the next section.

¹⁴⁶ For more discussion on positional goods see Sara J. Solnick & David Hemenway, *Is More Always Better?: A Survey on Positional Concerns*, 37 J. Econ. Behav. & Org. 373 (1998).

¹⁴⁷ Some argue that prioritization could improve overall network performance without degrading any user experience under certain conditions. See Ou, *supra* note 57. This would happen if, for example, the quality of websites on the fast lane increased without disturbing the performance of the websites on the slow lane. For example, if a group of emails and a time-sensitive finance application enter the network at the same time, they could block each other at the router. If an ISP were to prioritize the finance application and ensure that packets from that application are delivered in a timely manner to the users, that application clearly benefits. If the ISP slots the email packets between the finance application updates, the emails will be delivered in a timely manner without being dropped. This would ensure enhanced delivery of both the content in the slow lane and that in the fast lane. Ou argues that the relevant comparison should be between the user’s experience of the applications on the single-stream Internet versus those same applications on the slow lane of the bifurcated Internet—not the comparison between the application on the fast lane versus the application on the slow lane. *Id.*

¹⁴⁸ Riley & Topolki, *supra* note 137, at 2.

¹⁴⁹ ISPs would however take into consideration the impact of their prioritization procedures to the extent that loss of content providers affects the willingness-to-pay of their end-subscribers, but not the impact on the value of the Internet as a whole.

¹⁵⁰ See Litan, *supra* note 141.

¹⁵¹ 74 Fed. Reg. at 62,646, 62,651–52.

¹⁵² For example, the American Telemedicine Association refused to take a position on previous congressional efforts to mandate net neutrality. See AMERICAN TELEMEDICINE ASS’N, NET NEUTRALITY: THE CURRENT NEUTRALITY DEBATE AND ITS COMPETING PERSPECTIVES, AMERICAN TELEMEDICINE ASSOCIATION (2006), http://www.americantelemed.org/files/public/policy/Net_Neutrality.pdf.

¹⁵³ Letter from Jared Kopf, Chairman/President AdRoll.com, et al. to Julius Genachowski, Chairman, FCC (Oct. 19, 2009), *available at* <http://www.openinternetcoalition.org/index.cfm?objectID=69276766-1D09-317F-BBF53036A246B403>; see also Cecilia King, *Tech Venture Capitalists Lend Support to Net Neutrality Rules*, WASH. POST, Oct. 20, 2009, http://voices.washingtonpost.com/posttech/2009/10/tech_venture_capitalists_join.html (describing venture capitalists’ letter).

¹⁵⁴ Robin S. Lee & Tim Wu, *Subsidizing Creativity through Network Design: Zero-Pricing and Net Neutrality*, 23 J. ECON. PERSP. 73-74 (2009).

¹⁵⁵ 74 Fed. Reg. at 62,661.

